Time Evolution and Deterministic Optimisation of Correlator Product States

Vid Stojevic[1], Philip Crowley[1], Tanja Đurić [2], Callum Grey [1] Andrew Green[1]
[1] *London Centre for Nanotechnology, 17-19 Gordon St, London, WC1H 0AH*
[2] *Instytut Fizyki im. M. Smoluchowskiego, Uniwersytet Jagielloński, Łojasiewicza 11, 30-348 Kraków, Poland*

We study a restricted class of correlator product states (CPS) for a spin-half chain in which each spin is contained in just two overlapping plaquettes. This class is also a restriction upon matrix product states (MPS) with local dimension 2^n (n being the size of the overlapping regions of plaquettes) equal to the bond dimension. We investigate the trade-off between gains in efficiency due to this restriction against losses in fidelity. The time-dependent variational principle formulated for these states is numerically very stable. Moreover, it shows significant gains in efficiency compared to the naively related matrix product states - the evolution or optimisation scales as 2^{3n} for the correlator product states versus 2^{4n} for the unrestricted matrix product state. However, much of this advantage is offset by a significant reduction in fidelity. Correlator product states break the local Hilbert space symmetry by the explicit selection of a local basis. We investigate this dependence in detail and formulate the broad principles under which correlator product states may be a useful tool. In particular, we find that scaling with overlap/bond order may be more stable with correlator product states allowing a more efficient extraction of critical exponents - we present an example in which the use of correlator product states is several orders of magnitude quicker than matrix product states.

arXiv:1604.07210v1 [quant-ph] 25 Apr 2016

CONTENTS

I. Outline 2
II. Correlator Product States 3
III. TDVP with uCPS 5
IV. Properties of uCPS ground-state approximations 7
V. Quenches with uCPS 11
VI. Discussion 13
VII. Conclusions 14
Acknowledgments 15
A. Review of Matrix Product States and the Time-Dependent Variational Principle 15
B. Contracting uCPS with computational cost $O(D^3)$ 17
C. uCPS Gram matrix inverse pre-conditioning 17
References 18

Variational Ansaetze have underpinned many advances in the study of correlated quantum systems, BCS superconductivity and the Laughlin wavefunctions being notable examples. The essence is to identify key physical features of a quantum state such that they can be parametrised by relatively few numbers, providing an analytically or numerically tractable approximation. The idea of covering the system with coupled plaquette clusters is basis of several ansaetze. Examples include coupled cluster methods [1], contractor renormalization methods [2] and various hierarchical mean-field approaches [3]. Within all of these methods, the configurations and correlations of the single plaquette or cluster are evaluated exactly, however the inter-plaquette couplings are essentially treated in mean-field. Correlator product states (CPS - also referred to as entangled plaquette states) [4–13] go beyond this mean-field by fully parametrising the wavefunction on a set of overlapping plaquettes. Consistency between the wavefunctions in the regions of overlap communicates entanglement across the system.

The realisation of the central importance of the entanglement structure of quantum many-body states has led to the introduction of many variational families of tensor network states that efficiently parametrise it. The progenitor of these are matrix product states (MPS) [14, 15]. MPS have been extended to higher dimensions using projected entangled pair states (PEPS) [16], and to accommodate various aspects of the renormalisation of entanglement at critical points [17–19]. Such states naturally obey an 'area law', which, in any dimension, refers to the property that the entanglement entropy of a large enough region scales not with the volume, but with the area of the boundary of that region. It is proven for gapped systems in one dimension [20, 21] - and suspected in higher dimensions - that groundstates obey an area law and so can be efficiently described by tensor networks with relatively few components.

However, the ability to represent a state efficiently in no way implies that quantities such as expectation values of local operators can be calculated efficiently. Indeed, it has been demonstrated [22] that, unlike the computational complexity of contracting an arbitrary MPS, which is in P, the computational complexity of contracting a general PEPS network - the natural higher dimensional analogue - is #P-complete. In practice, except for very special tensor networks, in higher dimensions it is necessary to resort to either some approximate contraction scheme, or to statistical, Monte-Carlo type methods.

A possible strategy is to place additional restrictions upon the tensor network that render the quantities of interest easier to calculate. Such restrictions involve a compromise in accuracy, which must be balanced against gains in efficiency. Here, we consider the set of states at the intersection between MPS and CPS. The subset of CPS in which each spin is contained in only two plaquettes are also a subset of PEPS - the latter restricted to states with the same bond dimension and local Hilbert space dimension. Examples include stabiliser states

[23] and Kitaev's toric code [24]. In common with string bond states [25], of which they are a subset, CPS are efficiently sampleable, but not in general efficiently contractible, and can be used to calculate expectation values of local observables efficiently. The trade-off is that CPS are basis dependent - meaning that CPS defined with respect to different bases do not cover the same submanifold of the Hilbert space. Physical insights can motivate a judicious choice of basis for a given Hamiltonian.

We focus upon the above-mentioned subset of uniform CPS (uCPS), in one spatial dimension, and in the thermodynamic limit. We have two main aims:

i. *To develop the deterministic methods for optimising and evolving uCPS.* Our main tool will be a deterministic algorithm based upon the time-dependent variational principle (TDVP), an approach that projects the exact Schrödinger time-evolution onto the uCPS sub-manifold of the full Hilbert space. The deterministic TDVP approach is somewhat orthogonal to the usual statistical setting in which CPS are utilised.[1] This brings with it certain advantages, such as a very robust algorithm, as well as the ability obtain time-evolution information straightforwardly; but the main disadvantage is that with current techniques we are restricted to the one-dimensional setting. Nevertheless, many of the results provide insight about higher dimensional behaviour. Our construction will follow closely the TDVP implementation for MPS given in [28].

ii. *To provide a systematic study of basis dependence of uCPS and its interplay with computational speedup.* Speedup is expected due to the due to the ability to store and manipulate uCPS matrices more efficiently than their unrestricted uMPS counterparts, but this advantage is offset by the lower capacity of uCPS to encode ground state and quench information. Indeed, the interplay that we reveal is more subtle than one might anticipate.

I. OUTLINE

The basics of (u)CPS are reviewed in Section II, and their relation to (uniform) matrix product states (u)MPS is derived. The (u)CPS $\rightarrow$ (u)MPS mapping necessitates the introduction of the so called copier tensor. This construction is not sensitive to our one dimensional analysis, and it is the structure of the copier tensor in higher dimensions that is the starting point, and can provide indications, of how our one-dimensional results can be generalised.

In Section III the time-dependent variational principle (TDVP) for uCPS is derived. Since the cost of a single uMPS TDVP step scales as $O(dD^3)$, where d is the size of the physical and D of the virtual dimension, one naively expect the cost of each TDVP step to scale as $O(D^4)$, but we demonstrate that for uCPS the cost drops to $O(D^3)$.

Section IV provides a detailed study of the properties of uCPS ground state approximations, exemplified by a number of models: the Quantum Ising-, the Heisenberg-, and the XY-model (all spin-$\frac{1}{2}$). The imaginary time algorithm converges to the global minimum uniquely only for Hamiltonians that do not have degenerate groundstates. When the groundstate is degenerate it may still be possible to find a special choice of basis for which convergence is unique, however in general it turns out that uCPS breaks the degeneracy due to its basis dependence, and TDVP may converge to local minima associated with this breaking. When convergence is not unique, the algorithm acquires a probabilistic ingredient. For a Hamiltonian with a discretely degenerate groundstate the computational cost is increased only by a constant factor, but if the vacuum has a continuous degeneracy this factor seems to be larger than constant, and the number of possible local minima that TDVP can converges to seems to increase in an unbounded manner with bond dimension. Next, the efficiency of uCPS is compared with uMPS. We uncover an intricate picture. Naively one might expect that the uCPS in its optimal basis would capture a state with accuracy comparable to uMPS of which it is a restriction - *i.e.* uMPS with bond order equal to the local Hilbert space dimension. In fact, the situation is considerably worse; the plaquette overlap required to match the accuracy of a uMPS of a given bond dimension is proportional to that bond dimension. This is exponentially worse than naive expectation since the local Hilbert space dimension scales exponentially with the size of the overlap region. However, in the optimal basis, we find that the computer *time* needed to converge the ground state approximation to a desired accuracy scales in the same way for both uCPS and uMPS. We also study the scaling of quantities, such as energy, entanglement entropy, or the correlation length, with bond dimension in detail. The scaling behaviour is generally basis dependent, but much smoother for uCPS than for uMPS. This allows for a more accurate estimates of certain physical quantities to be made with uCPS than with uMPS with the same computational time cost (but, still, at exponentially larger computer memory cost). In general we find that an optimal basis corresponds to one that is aligned as closely as possible with the entanglement generating terms in the Hamiltonian.[2]

In Section V, real time TDVP is applied to uCPS in order to simulate quantum quenches. We concentrate in particular upon quenches across a critical point that exhibit dynamical phase transitions. It is found that the capacity of uCPS to correctly capture a quench to a desired accuracy scales, analogously to the static context above, exponentially worse than uMPS in memory requirements but in the same manner as far as time costs are concerned. This is again only true provided one is working in the optimal basis, or reasonably close to it; but unlike the static case, here there are no probabilistic aspects to contend with. It is also found that for a sub-optimal basis choice, the uCPS approximation to the quench can completely miss the dynamical phase transition. In fact, with the same basis choice, the uCPS approximation of the ground state misses the equilibrium phase transition. All of this gives some indication of the types of issues one may encounter with CPS in general, if working in a sub-optimal basis.

We elaborate upon these issues in Section VI and give a final brief summary in Section VII. To keep things as self-

[1] Although other deterministic methods have been developed in [26] and [27].

[2] This is clearly only a loose guiding principle, discussed at length in Section IV, and exemplifies the issue one needs to address for CPS/string bond states in general, of identifying the optimal basis for the physics one wishes to study.

contained as possible, a basic review of MPS and TDVP is provided in Appendix A. Appendix B details the contraction ordering that achieves the optimal $O(D^3)$ scaling of a single uCPS TDVP step, while details of a preconditioning step necessary for an iterative sub-routine in the TDVP algorithm to scale optimally is provided in Appendix C.

II. CORRELATOR PRODUCT STATES

In this section the basics of correlator product states (CPS) and the relationship of the CPS variational class to string bond states will be reviewed. We then focus upon the main subject of this paper, uniform correlator product states (uCPS) in the thermodynamic limit, detail how the uCPS class can be understood as a subset of uniform matrix product states (uMPS), and discuss a number of implications. For a brief review of Matrix Product States (MPS) the reader is referred to Appendix A.

Definition of CPS - A CPS wavefunction on an N-site spin lattice is formed by dividing the lattice into a set of overlapping subsets called plaquettes - labelled here by P. Wavefunctions are defined over each plaquette by a tensor C_P and a correlated product between them taken according to

$$|\Psi[C]\rangle = \sum_{i_1,i_2\cdots i_N} \prod_P C_P^{i_1,\cdots i_l} |i_1, i_2, \cdots, i_N\rangle \qquad (1)$$

in order to maintain consistency in the regions of overlap and thereby entangling spins in disjoint plaquettes. The plaquette sizes are taken to be independent of the size of the system, and sufficiently small that the wavefunction on each can be stored and manipulated with small computational cost. The accuracy of this Ansatz is influenced by a large number of factors, including the structure of the Hamiltonian, the choice of basis, and size and number of overlaps between the plaquettes.

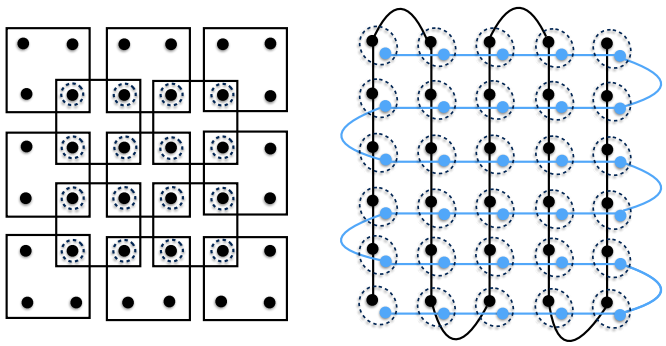

FIG. 1. (Colour online). The diagram on the left hand side depicts a CPS in two dimensions with single-site overlaps. The squares represent the C tensors, and the black dots denote the spins, each C tensor thus having four indices. The right hand side depicts a string bond state consisting of a product of two matrix product states, each covering the entire lattice. The dashed circles are present at sites where the underlying CPS/string bond state tensors overlap, and each such circle can be understood as representing the copier tensor projecting the indices in the overlap onto the physical spin.

The key requirement for the Ansatz of the form (1) to be amenable to Monte Carlo type algorithms is that the coefficients of the wavefunction $|\Psi[C]\rangle$ in (1) be efficiently calculable.[3] This is clearly satisfied whenever the coefficients C_P are efficiently calculable individually. For CPS this requirement holds simply because the plaquettes are by construction taken to be of a small constant size, but more generally *any* efficient format for C_P will be a priori equally suitable. The definition of the class of *string bond states* [25] is also of the general form (1), but with the sub-lattices P taken to be strings of arbitrary length instead of plaquettes, and the C_P coefficients given by a matrix-product form. Examples of both a CPS with single-site overlap, and a string bond state given by a product of two MPSs, each covering the entire lattice, are depicted in Figure 1.

Since a general-form wavefunction on a constant size plaquette can always be recast in MPS form, CPS clearly form a subset of string bond states. It is, however, not the case that in general either CPS or String Bonds States can be efficiently represented as PEPS (or MPS in one dimension). The number of overlapping plaquettes for a CPS, or strings for a string bond state, is allowed to be of the order of system size, yet representing such a state as PEPS/MPS requires in general an exponentially large bond dimension. The Laughlin state, for example, can be written as a 2-site CPS [4], but the number of overlapping CPS tensors at each site is equal to the system size; a diagram depicting the general form of the Laughlin state is given in Figure 2.

A one dimensional CPS with open boundary conditions, and with the restriction that only two plaquettes overlap, is given by:

$$|\Psi[C]\rangle = \qquad (2)$$
$$\sum_{i_1,i_2\cdots i_N} v_L^{i_1} C_1^{i_1 i_2} C_2^{i_2 i_3} \cdots C_{N-1}^{i_{N-1} i_N} v_R^{i_N} |i_1, i_2, \cdots, i_N\rangle \ .$$

The copier tensor - In order to write down a tensor network for a CPS it is necessary to introduce the copier tensor. For the state (2), for example, the copier is given by $\mathbb{1}_{ijk}$, defined as the tensor with components equal to unity when all its indices are equal, and zero otherwise. It will be convenient to represent the copier graphically as $\curlyvee$; the CPS state in (2) can then be represented as:

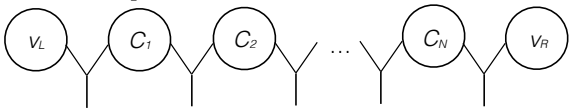

,

where any variation of the dimensionality of the indices in (2) along the chain can be indicated by an additional label if necessary. As each physical index in (2) corresponds to the overlap of the corresponding CPS plaquette, for an n-site overlap the dimensionality of the index is obtained by grouping the n fundamental spins s together, and is given by s^n. The dimensionality of these indices therefore grows exponentially with the size of the overlap. The basis dependence of CPS/string bond states is made explicit by the non-covariant definition of the copier tensor, and it is easily verified that a rotation by a local unitary on the physical leg can not in general be "pulled through" to the virtual level.[4]

---

[3] This is only a necessary condition, and is not sufficient to guarantee that the Monte Carlo algorithm will converge.

[4] One can consider adding such local unitaries, or more general matrices, in order to enlarge the CPS variational class and restore rotation independence.

The copier tensor generalises straightforwardly to more general CPS, and to higher dimensions. For example, the dashed circles in Figure 1 should in fact be understood as representing three-index copier tensors. For the Laughlin state depicted in Figure 2, the dashed circles again represent copiers, but the number of virtual legs is now of the order of system size.

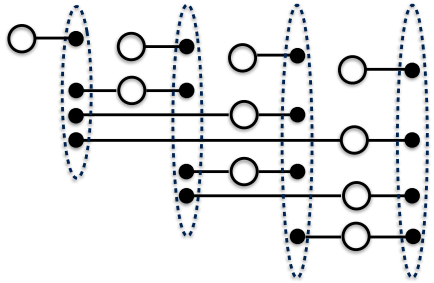

FIG. 2. This diagram depicts the general form of a tensor network corresponding to the CPS representation of the Laughlin state over four spins, and exemplifies a CPS for which the number of overlapping CPS tensors at each physical spin is of the order of system size. The circles denote the CPS tensors, with their indices depicted by lines ending in the black dots, and the dashed ovals represent the action of the copier tensors mapping overlaps of the CPS tensors onto the four physical spins. The diagram corresponding to the CPS representation of a Laughlin state over an arbitrary number of spins follows the same pattern.

Furthermore, using the tensor network representation of CPS, a natural generalisation of the CPS/string bond state variational class can be achieved by promoting the copier to any tensor that is sufficiently sparse to leave sample-ability intact. As an example, one can imagine a CPS-type Ansatz for which the copier is replaced by a tensor the non-zero components of which are given by an n-site MPS.

Mapping one-dimensional CPS into MPS - One dimensional CPS of the form (2) form a sub-class of MPS with open boundary conditions,

$$|\Psi[A]\rangle = \sum_{i_1,i_2,\cdots,i_N}^{d} v_L^\dagger A_1^{i_1} A_2^{i_2} \cdots A_N^{i_N} v_R \, |i_1, i_2, \cdots, i_N\rangle \ , \tag{3}$$

provided that the physical spin- and bond-dimensions of the MPS tensors, which at site K we denote as $d^{(K)}$ and $D^{(K)}$ respectively, are the same either to the left or to the right (i.e. provided that either $d^{(K)} = D^{(K-1)}$, or $d^{(K)} = D^{(K)}$). There are many ways to map the CPS defined in (2) into an MPS of this type. One possibility is the following:

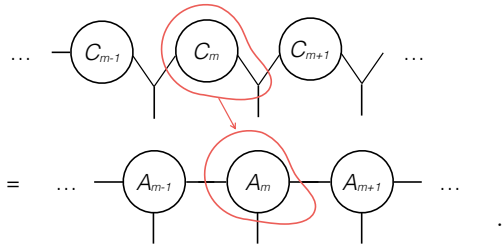

Alternatively the C matrix can be placed to the right of the copier in the identification, or placed both to the left and the right of the copier after decomposing each C matrix into a product of two matrices.

In higher dimension, any CPS/string bond state with a constant number of overlapping plaquettes/strings can be mapped to a PEPS with small bond dimension using analogous considerations.

Uniform Correlator Product States (uCPS) - This paper is concerned mainly with uniform translation invariant CPS in the thermodynamic limit. In order to achieve a fully translation invariant representation, it is necessary that the plaquettes contain an even number of spins, so that half of the spins of each plaquette overlap with half of the spins of a neighbouring plaquette.[5] This means that the C matrices are square and that the bond dimension along the chain is constant. These states will be referred to as uniform correlator product states (uCPS), and form a proper subset of the uMPS class described in Appendix A. Graphically a uCPS is represented as:

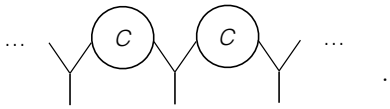

The uCPS $\rightarrow$ uMPS mapping is achieved in the same manner as in the non translation invariant setting described above, with analogous freedom in how this mapping is realised.

The increase in uCPS bond dimension is achieved by increasing the size of the overlap between plaquettes. If n is the number of fundamental spins in each overlap, for a lattice of s-level spins the size of the C matrices is $s^n \times s^n$, and the bond dimension D, given by $D = s^n$, therefore increases exponentially with n.

The uCPS transfer matrix - The transfer matrix $\tilde{E}$ (see equation (A6)) corresponding to the uCPS $\rightarrow$ uMPS mapping with the C matrix placed to the right of the copier tensor takes the form:

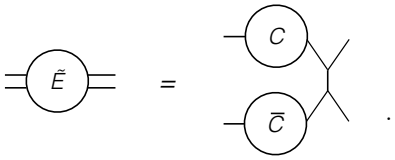

Due to the sparseness of the copier tensor, $\tilde{E}$ has only D non-zero eigenvalues. This is related to the fact that the the object is a projector that effectively reduces the size of a D^2 dimensional index to a D dimensional one. It also follows from this that the uCPS variational class is in fact a subclass of reduced rank - not full rank - uMPS. It is sufficient therefore to work with the projected version of the above object, E:

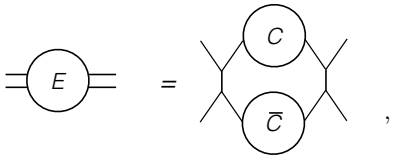

the non-zero subspace of which is given by the $D \times D$ dimensional matrix:

$$E_{ij} = C_{ij}(.*)\overline{C}_{ij} \ . \tag{4}$$

[5] More generally one can consider representation that are invariant under translations by m sites, with $m > 1$.

The operator $(.*)$ denotes component-wise multiplication for repeated indices. As for uMPS, in order to ensure finite normalisation of the state, it is necessary to normalise the C matrices so that the largest eigenvalue of E is one.

As discussed in Appendix A, the left and right eigenvectors of the transfer matrix $\tilde{E}$ corresponding to eigenvalue one are centrally important objects, which determine the Schmidt coefficients across a cut of the infinite chain into two semi-infinite intervals, and, more generally, represent the cumulative effect of the environment when calculating the expectation value of a local operator, from its insertion to $\pm\infty$. In order to fully utilise the special form of uCPS over a generic uMPS in the TDVP algorithm (see Section III), it is necessary to work directly with the D dimensional left and right eigenvectors of E (4) corresponding to eigenvalue one (rather than with the eigenvectors of $\tilde{E}$), which we will refer to as V_L and V_R, respectively (these should not be confused with the boundary v_L and v_R of open-boundary finite MPS (3)). We note that, while the left eigenvector of $\tilde{E}$ is given by:

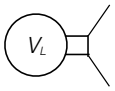

, (5)

the right eigenvector explicitly depends on C as:

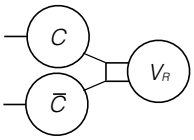

. (6)

Had we made the identification with uMPS by placing the C matrices to the right of the copier tensor, the situation would be the reverse of the above, with the right eigenvector depending only on V_R, and the left on V_L, C, and $\overline{C}$.

A uMPS can always be gauge-transformed (A8) to the left or right canonical gauge, as discussed in Appendix A, meaning that either the left or right environment matrix is equal to the unit matrix. This is in general not achievable with uCPS, since the set of uMPS gauge transformations that preserve the uCPS form is extremely restricted, and the remaining gauge freedom is not sufficient in general to achieve the left or right canonical gauge. Stated in another manner, uCPS corresponds to a very restricted type of uMPS, and imposing such a restriction on uMPS fixes nearly all of its gauge freedom. One of the implications of all of this is that it does not seem possible, at least straightforwardly, to implement an analogue of the iDMRG algorithm [29] for uCPS, since it relies upon the ability to achieve the left and right canonical gauges and to switch between them. It is interesting that there are nevertheless two natural gauge choices for uCPS, as explained above, related to the freedom of placing the copier to the left or to the right of the C matrix in the uCPS $\rightarrow$ uMPS mapping.

Computational cost of contracting uCPS - Since a uCPS maps to a uMPS with equal bond and physical dimensions, i.e. $d = D$, and the cost of calculating a local uMPS observable is $O(dD^3)$ (see Appendix A), one naively expects that the cost of calculating expectation values of local observables for uCPS scales as $O(D^4)$. It turns out, however, that one can do better. One possible reduction in cost occurs due to the fact that the copier tensor factorises, and that a uCPS with overlap of size n and local spin of dimensionality s, can be mapped to an n-site[6] uMPS with bond dimension $D = s^n$. How this is achieved is most clearly demonstrated graphically. For example, for $n = 3$ and $D = s^3$ the uMPS tensor of the uCPS form can be decomposed into 3 MPS tensors as:

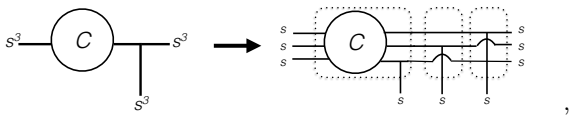

,

where the dashed regions identify the three MPS tensors with physical dimension s.[7] The cost of calculating local expectation values for an n-site uMPS can therefore be reduced to $O(\log(D)sD^3) = O(\log(D)D^3)$, simply by applying an optimal contraction ordering for n-site uMPS after making the above decomposition.

The presence of the $D^2 \rightarrow D$ projector χ in the transfer matrix $\tilde{E}$ hints that an even lower cost may be achievable. This turns out to be the case, and in the next section we demonstrate that, in addition to the computation of expectation values of local operators, all the steps of the uCPS TDVP algorithm can be computed with cost $O(D^3)$.

III. TDVP WITH uCPS

The time-dependent Schrödinger equation, $\frac{d}{dt}|\Psi(t)\rangle = -i\hat{H}(t)|\Psi(t)\rangle$, does not in general have an exact solution if the state vector is restricted to remain within some class of variational states $|\Psi(V_X)\rangle$ during the evolution, where V_X denotes some set of variational parameters labelled by an index X. This can be observed by considering the equation:

$$\sum_X \dot{V}_X \frac{\partial}{\partial V_X}|\Psi(V_X)\rangle + i\hat{H}(t)|\Psi(t)\rangle = 0 \ , \qquad (7)$$

where the second term can correspond to an arbitrary direction in Hilbert space, but the first is highly restricted. Instead, one must adopt some procedure for obtaining an optimal $\dot{V}_X$ that minimises the left hand side of (7) with respect to some cost function. The time-dependent variational principle (TDVP) corresponds to the natural cost function in a quantum mechanical setting, namely the 2-norm, and the solution to $\dot{V}_X^*$ is given by: [8]

$$\dot{V}_X^* = \arg\min_{\dot{V}_X} \left\| \sum_X \dot{V}_X \frac{\partial}{\partial V_X}|\Psi(V_X)\rangle + i\hat{H}(t)|\Psi(t)\rangle \right\|^2 \ . \qquad (8)$$

The full variational space of uCPS consists of a general variation dC of the uCPS matrix C. In the context of the uCPS

---

[6] A n-site uMPS is a MPS in the thermodynamic limit such that its tensors are unchanged after a translation by n sites.

[7] Clearly all the uCPS information has been incorporated into the left-most tensor for this particular choice of decomposition, but if desired C can be distributed over the n MPS tensors in a symmetrical manner by performing a suitable factorisation.

[8] See [28, 30] and Appendix A in reference to the application of TDVP to (u)MPS.

→ uMPS mapping described in Section II this corresponds to a restricted uMPS variation of the form:

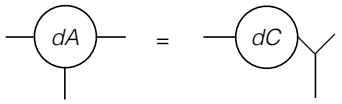

(9)

and can be understood as a tangent vector to the uCPS manifold (see also Appendix A). Writing this out symbolically:

$$|\Phi(dC, C)\rangle := \sum_{i,j} dC_{ij} \frac{\partial}{\partial C_{ij}} |\Psi(C)\rangle \equiv dC^{ij} |\partial_{ij}\Psi(C)\rangle \ , \quad (10)$$

the solution to the minimisation problem (8) is given by:

$$\sum_{k,l} G_{\overline{ij}|kl} \dot{C}^{kl} = -\, i \left\langle \partial_{\overline{ij}}\Psi(C(t)) \right| \hat{H}(t) \left|\Psi(C(t))\right\rangle \ , \quad (11)$$

where G is the uCPS Gram matrix (c.f. (A10)):

$$G_{\overline{ij}|kl} := \left\langle \partial_{\overline{ij}}\Psi(C) \right| \partial_{kl}\Psi(C)\rangle \ . \quad (12)$$

However, the solution in (11) contains terms that diverge if the variations dC are left fully unconstrained. The full tangent space spanned by $|\Phi(dC, C)\rangle$ includes transformations along the state $|\Psi(C)\rangle$ itself, and it is precisely such norm-changing variations that generate divergences. These must be projected out, which can be achieved by implementing the uCPS analogue of the uMPS left or right tangent space gauge condition (A9). Moreover, the tangent space gauge condition simplifies the form of the Gram matrix, which becomes local, meaning that all the contributions to (12) vanish except those for which the variations dC in the bra and ket occur at the same lattice site.

Gauge fixing for uCPS - For the sake of concreteness we will demonstrate how to impose the *left* tangent gauge condition for uCPS, which reads:

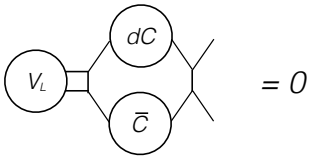

or symbolically

$$V_{(L)i}(.*)dC_{ij}(.*)\overline{C}_{ij} = 0 \ . \quad (13)$$

A solution to this equation is given by dC of the form:

$$dC_{ij} = \sum_{\tilde{\alpha}\beta} B_{\tilde{\alpha}\beta} V^{\tilde{\alpha}}_{(R)i} V^{\beta}_{(L)j}(.*)(1(./)\overline{C}_{ij}) \ , \quad (14)$$

where B is an arbitrary $(D-1)\times D$ matrix, $(./)$ denotes component-wise division, and $V^{\alpha}_{(L)}$ and $V^{\alpha}_{(R)}$ are the D left and right eigenvectors of the transfer matrix (4), so that

$$E = \sum_{\alpha} \lambda_\alpha V^{\alpha}_{(L)} V^{\alpha}_{(R)} \ . \quad (15)$$

By convention $V^1_{(L)} \equiv V_{(L)}$, and $V^1_{(R)} \equiv V_{(R)}$, and the $\tilde{\alpha}$ index in (14) thus runs from 2 to D.

With the tangent space gauge condition in place, the analogue of (11) for B is given by:

$$\sum_{\tilde{\gamma}\delta} G_{\overline{\tilde{\alpha}\beta}|\tilde{\gamma}\delta} B^{\tilde{\gamma}\delta} = -\, i \left\langle \partial_{\overline{\tilde{\alpha}\beta}}\Psi(C(t)) \right| \hat{H}(t) \left|\Psi(C(t))\right\rangle \ , \quad (16)$$

where $\partial_{\overline{\tilde{\alpha}\beta}}$ on the right hand side stands for a derivative with respect to $\overline{B}^{\tilde{\alpha}\beta}$, after making use of the chain rule applied to a differential dC of the form (14). After deriving the solution for $B_{\tilde{\alpha}\beta}$, by applying the inverse of the Gram matrix to (16), the TDVP update for C is obtained via (14).

Efficient Implementation - As discussed at the end of Section II, naively following an optimal uMPS contraction ordering in the uCPS setting will yield a computational cost that scales at best as $O(\log(D)D^3)$ for a local Hamiltonian, and we expect to be able to do better. Details of a contraction ordering such that the cost of calculating the right hand side of (16) scales as $O(D^3)$ are given in Appendix B.

In addition to this, it is necessary to find an $O(D^3)$ implementation for the action by the inverse of the $D^2 \times D^2$ matrix G on equation (16). Stemming from the fact that our parameterisation of variations dC satisfying the tangent left gauge condition (14) depends in a complicated way on the eigenvalue decomposition of the transfer matrix, the uCPS Gram-matrix turns out to be a much more complex object than its uMPS equivalent (given simply by $\rho_l \otimes \rho_r$ (A13)), and is given by:

$$G^{\tilde{\gamma}\delta|\tilde{\alpha}\beta} = \sum_{ij} \lambda_{ij} T^{\tilde{\gamma}\delta|\tilde{\alpha}\beta}_{ij} \ , \quad (17)$$

where

$$T^{\tilde{\gamma}\delta|\tilde{\alpha}\beta}_{ij} = \left(\overline{V}^{\tilde{\gamma}}_{(R)\ i} \overline{V}^{\delta}_{(L)\ j}\right)(.*)\left(V^{\tilde{\alpha}}_{(R)\ i} V^{\beta}_{(L)\ j}\right) \ , \quad (18)$$

and

$$\lambda_{ij} = V^0_{(L)\ i} V^0_{(R)\ j}(.*)(1(./)E_{ij}) \ . \quad (19)$$

Naively, both the cost of calculating G^{-1}, as well as its action on a vector, scales as $O(D^6)$; an obvious approach to reduce this cost would be to attempt to write G in a decomposition that would allow the action of its inverse on a vector to be calculated with cost $O(D^3)$, analogous to the manner in which a $O(d^2D^6) \to O(dD^3)$ reduction in cost is achieved for uMPS TDVP. As far as we have been able to ascertain, this is not possible for (17). While such an inverse *does* exist for the matrix obtained by replacing all tilde-indices in (17) by their non-tilde extensions, and is given by inverting all the constituent matrices and taking $\lambda_{ij} \to 1(./)\lambda_{ij}$, the truncation of the eigenvalue-one eigenvector seems to conclusively obstruct achieving any appropriate decomposition for G^{-1} itself. It is nevertheless possible to use an iterative method (such as the biconjugate gradient algorithm) in order to calculate the action of G^{-1} and retain $O(D^3)$ efficiency, since this only requires that the action of G on an arbitrary vector be calculable with cost $O(D^3)$.

A further caveat to the above is that in order to achieve $O(D^3)$ scaling, the number of iterations in the iterative subroutine required to achieve some desired accuracy must scale as a constant for large enough D, and this is not guaranteed. For example, the number of iterations *does* scale worse than constant if the pre-conditioning step is not implemented appropriately (its implementation is described in Appendix C). However, beyond this observation, for all the examples studied in this paper the iterative subroutine is observed to scale as $O(D^3)$.

Finally, it should be noted that while the general uMPS TDVP algorithm does not require an iterative subroutine for the Gram-matrix inverse step in order to achieve optimal efficiency, it does require such a subroutine in order to calculate

the third term in (A14) with cost $O(dD^3)$. The same term in uCPS TDVP, on the other hand, can be calculated explicitly with cost $O(D^3)$, as is described in Appendix B.

IV. PROPERTIES OF uCPS GROUND-STATE APPROXIMATIONS

In this section the imaginary time uCPS TDVP algorithm is used to obtain uCPS ground state approximations for a number of exemplary models: the quantum Ising model in a transverse magnetic field, the XY-model, and the Heisenberg model (all spin-$\frac{1}{2}$). The general aim is to study the capacity of uCPS to capture ground state properties both at and away from criticality, while carefully considering the effects of uCPS basis-dependence. We provide a detailed study of the convergence properties of uCPS ground state energy estimates, and of the convergence/scaling[9] behaviour of the entanglement entropy and of the correlation length, all with respect to the size of the uCPS overlap n (or, equivalently, with respect to the bond dimension D, related to the overlap by $D = s^n$). In addition, the performance of the uCPS TDVP algorithm is analysed, and compared with the general uMPS imaginary time TDVP algorithm applied to the same models.

Basis dependence - In order to demonstrate the effect of uCPS basis choice, we calculate the uCPS groundstate energy estimates for the transverse field Ising Hamiltonian:

$$\hat{H} = \sum_{i\in\mathbb{Z}} -J\hat{\sigma}_i^z\hat{\sigma}_{i+1}^z + h\hat{\sigma}_i^x \ , \tag{20}$$

the orientation of which remains fixed while the uCPS basis is rotated from the x- to the z-direction. Here $\{\hat{\sigma}^x, \hat{\sigma}^y, \hat{\sigma}^z\}$ are the Pauli matrices, J determines the coupling strength between nearest neighbour spins, and h determines the strength of the magnetic field; the model is critical for $\frac{h}{J} = \pm 1$.

The results for 2-site uCPS overlap ($D = 4$), at $h = 1$, are displayed in the plot in Figure 3. At the points at which the two branches cross, imaginary time TDVP converges to a unique global minimum; this occurs for rotations away from the z-basis by integer multiples of $\frac{\pi}{2}$, i.e. for z- and x- basis choices, and at one additional intermediate angle, the precise value of which depends upon uCPS overlap. Away from the special points at which the branches cross in Figure 3, TDVP converges to the global minimum only for a certain fraction of TDVP runs initiated from a random uCPS state.

The branching occurs due to the fact that, except for very special basis choices, uCPS breaks the two-fold degeneracy that is present in the uMPS approximation of the ground state at $h = 1$.[10] Thus, two states that would have exactly the same energy in the uMPS approximation acquire slightly different energies with uCPS due to the rotation dependence of the uCPS Ansatz. TDVP converges to one of these two solutions, depending upon details of the initial random state. At present, we do not understand precisely how the outcome is encoded in the initial state. A hint might perhaps be contained in the fact that a local unitary cannot be pushed through the copier tensor, so that the CPS parametrisation divides the Hilbert space into topologically distinct classes [32]. As one would expect, the branching is present for all values of h in the ferromagnetic phase, and the two branches collapse as one enters the paramagnetic phase.

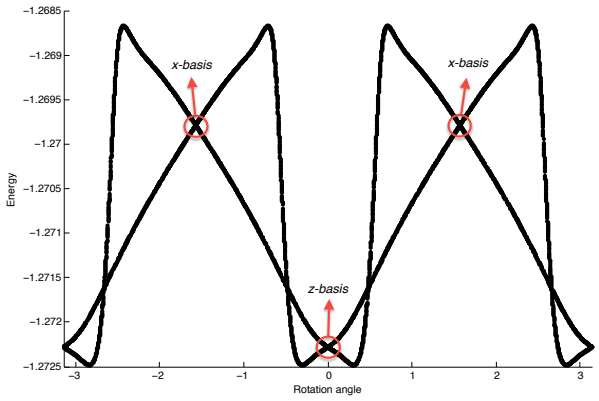

FIG. 3. (Colour online) The behaviour of 2-site overlap uCPS ground state energy approximations for the quantum Ising model at criticality, as a function of the global rotation angle around the y-axis away from the z-basis choice. The TDVP algorithm exhibits unique convergence to the global minimum only at the points marked by the red circles, which denote the z- and x- basis choices, and at one intermediate point. In general, depending upon the details of the initial random state, imaginary time TDVP will converge to one of the minima associated with the splitting of otherwise degenerate energy levels, which occurs due to the lack of rotation invariance of the uCPS Ansatz.

In general, imposing constraints on uMPS can potentially introduce local minima on the variational manifold. For example, uMPS TDVP often gets stuck in local minima when uMPS matrices are restricted to be real, and in that case some of the local minima can be singular.[11] In the present situation the local minima do *not* correspond to singular points on the uCPS manifold, and thus no associated numerical complications are encountered. Moreover, as long as the number of branches is constant as a function of uCPS overlap, in order to be certain of having reached the ground state with some probability, the algorithm needs to be run a fixed number of times: $O(2)$ times for the example of the quantum Ising model in the ferromagnetic branch away from the special points.

Let us next consider the XY-model, which is described by the Hamiltonian:

$$\hat{H} = -\sum_{i\in\mathbb{Z}} \frac{1+\gamma}{2}\hat{\sigma}_i^x\hat{\sigma}_{i+1}^x + \frac{1-\gamma}{2}\hat{\sigma}_i^y\hat{\sigma}_{i+1}^y + h\hat{\sigma}^z \ . \tag{21}$$

[9] The entanglement entropy and the correlation length do not converge at criticality, so only the scaling properties of these quantities can be considered.

[10] It should be noted that the *exact* ground state of the quantum Ising model is doubly degenerate for $h < 1$, but at finite bond dimension the $h = 1$ uCPS/uMPS approximations are still effectively in the ferromagnetic phase (see e.g. [31], and also the discussion relating to Figure 11).

[11] Meaning that TDVP has in fact converged to a uMPS of a lower bond dimension than that of the initial random state. As some eigenvalues of the Gram-matrix go to zero in this limit, the TDVP algorithm becomes numerically unstable, at least in its naive implementation (for an implementation of MPS TDVP that avoids such instabilities see [33]).

The ground state is two-fold degenerate for $0 < h < 1$ and $0 < \gamma < 1$. In this regime we observe two branches, as in the case of the quantum Ising model. For $\gamma = 0$ the symmetry of the Hamiltonian is increased to $U(1)$, and here the number of local minima that uCPS converges to is not bounded as the size of the uCPS overlap is increased, as far as our numerical study is able to ascertain, and the $O(D^4) \to O(D^3)$ gains in the scaling of computational costs associated with the special structure of uCPS seem to be lost for this model. In addition, no basis choice exists for which TDVP always converges to a global minimum.

When the symmetry is increased to $SU(2)$, however, as e.g. for the spin-1/2 Heisenberg model, described by the Hamiltonian:

$$\hat{H} = \sum_{i \in \mathbb{Z}} J \left(\hat{\sigma}_i^x \hat{\sigma}_{i+1}^x + \hat{\sigma}_i^y \hat{\sigma}_{i+1}^y + \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z \right) \ , \qquad (22)$$

multiple branches are no longer observed. Full rotation invariance means that optimal uCPS approximations of physical quantities must be completely independent of the choice of basis, and moreover, in the case of the Heisenberg model, TDVP is always observed to converge to a global minimum.

At this stage it can be observed that, in a loose sense, a judicious basis choice corresponds to one that is 'optimally' aligned with the entanglement generating terms in the Hamiltonian. Clearly, for the quantum Ising model at criticality, the z-basis is maximally aligned with $\hat{\sigma}^z \otimes \hat{\sigma}^z$, and as demonstrated provides better energy estimates than those obtained by picking the x-basis at equal overlap. Another desirable property with this choice is that TDVP always converges to the global minimum. However, as indicated by the plot in Figure 3, the best energy estimate is in fact achieved at an intermediate angle, but the computational downside of picking this point is that TDVP may get stuck in a local minimum, and that one needs to scan in order to find the orientation that achieves the lowest energy. In addition, the behaviour of uCPS is clearly highly model dependent. For example, the XY-model for $\gamma = 0$ has *no* basis choice at which convergence is unique, so scanning in this context may be a more sensible strategy than for the quantum Ising model. Moreover, due to the $U(1)$ invariance of this model, the number of minima that TDVP converges to seems to increase indefinitely as overlap size is increased. Since the TDVP algorithm converges to the global minimum only for a small fraction of runs, the algorithm becomes essentially statistical, and it is therefore questionable whether imaginary time TDVP has any advantages here over Monte Carlo, the standard approach for CPS/String Bond type Ansaetze. For the $SU(2)$ invariant Heisenberg model the symmetry is large enough to eliminate such issues, but since all basis choices yield the same approximations for physical quantities, there is little motivation left for using a basis dependent Ansatz in the first place. In conclusion, what is meant by an 'optimal basis' is a function of both the model under investigation, and what it is that one wishes to achieve. uCPS provides a good demonstration, in a well controlled context, why identifying a good basis is such a difficult problem for general CPS/string bond states. Nevertheless, attempting to maximally align the basis with entanglement generating terms certainly seems to be a good general guiding principle.

Ground state convergence at criticality - Next we investigate the ground state convergence properties of the transverse field quantum Ising model at criticality, comparing uCPS in the z- and x- bases with uMPS. As noted, TDVP always converges to the global minimum for these basis choices. The plot in Figure 4 demonstrates that the uCPS ground state energy estimates are better for the z- than for the x-basis choice (that this is the case for 2-site uCPS overlap is already evident from Figure 3). The more surprising observation is that, while the uMPS energy converges roughly polynomially as a function of bond dimension $D_{(\mathrm{uMPS})}$ (since the model is critical), the uCPS energy converges instead polynomially *with the size of the uCPS overlap* - note that the x-axis in the plot is shared between uCPS overlap size n and $D_{(\mathrm{uMPS})}$. Thus, since $D_{(\mathrm{uCPS})} = 2^n$, to achieve the same order of accuracy in the energy estimates uCPS requires an exponentially larger bond dimension than uMPS. Complementing this result is another surprising observation, depicted in the plot in Figure 5: the total time needed to reach convergence to same accuracy scales in the same way for uMPS with respect to bond dimension as for uCPS with respect to overlap size.

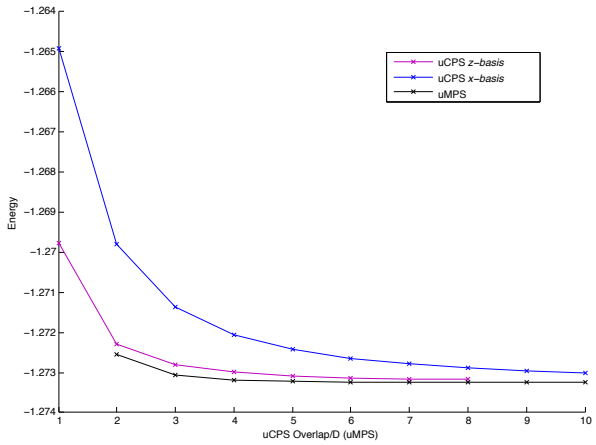

FIG. 4. (Colour online). Convergence of uCPS and uMPS ground state energies for the transverse Quantum Ising model at the criticality. The x-axis denotes overlap size for uCPS data, and bond dimension for uMPS data.

Thus, the indications from this example are that, even though the *computer memory* costs are exponentially larger for uCPS than for uMPS, the *computer time* needed to reach some desired accuracy for the ground state energy approximation scales in the same way for both uCPS and uMPS. This result seems to be generic, and is observed to hold for all the models studied in this paper.

In this context it is also useful to consider the $SU(2)$ invariant Heisenberg model, so that it is impossible to pick out some preferred basis choice. While one avoids the complications associated with basis dependence, the result of using an orientation dependent Ansatz on a rotation invariant Hamiltonian seems to manifest itself in extremely slow convergence, as a function of uCPS overlap, of the ground state energy approximations compared to what is observed for models whose entanglement generating terms point in a definite direction, such as e.g. the quantum Ising model. This is illustrated by the plot in Figure 6. In addition, the times required for TDVP to converge are much longer, at equal uCPS overlap, than for the quantum Ising model. For example, the maximal bond dimension that we could achieve in reasonable time is $D_{\mathrm{uCPS}} = 2^6$, and the corresponding energy density reproduces the correct result only to three digits. This slowdown originates in the iterative subroutine of the TDVP algorithm,

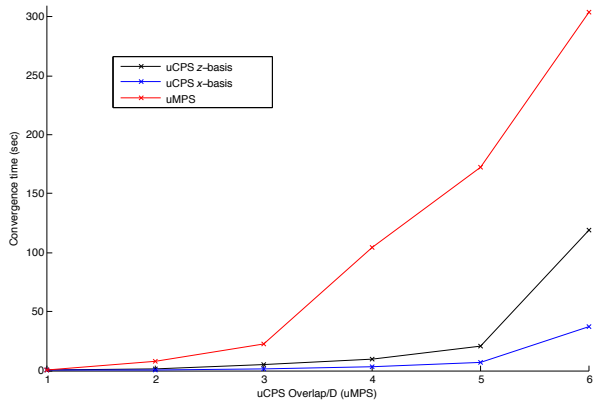

FIG. 5. (Colour online). Convergence times for uCPS and uMPS approximations for the ground state of the quantum Ising model at criticality. The x-axis denotes overlap size for uCPS data, and bond dimension for uMPS data.

which requires a much larger number of iterations to converge to the required accuracy.

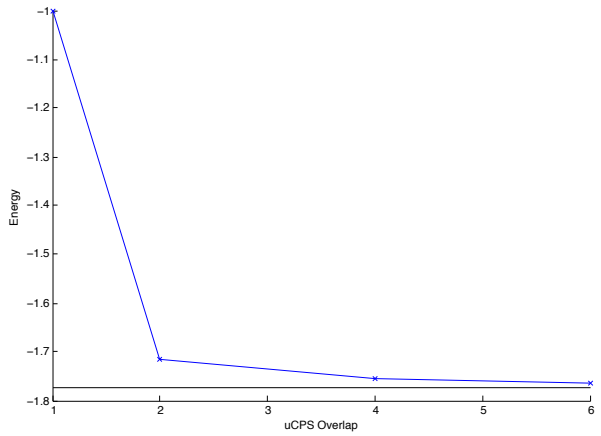

FIG. 6. (Colour online). Convergence of uCPS ground state energies for the Heisenberg model. The black horizontal line denotes the exact energy. Whilst the uCPS approximations are not basis dependent, convergence turns out to be very slow.

Ground state convergence in a gapped phase - Next we examine the ground state convergence of uCPS deep in the gapped phase. The observations made above, regarding the computational time and memory costs of uCPS vs. that of uMPS, remain true away from criticality. Here we shall demonstrate that the analogous observations also hold for two rather different types of physical quantities, namely the *correlation length* μ and the *entanglement entropy* S. The correlation length is obtained from the largest eigenvalue λ_2 of the transfer matrix smaller than 1 as:

$$\mu_{\mathrm{uMPS}} = -\frac{1}{\log(\lambda_2^{(\mathrm{uMPS})})} \quad , \quad \mu_{\mathrm{uCPS}}(n) = -\frac{n}{\log(\lambda_2^{(uCPS)})} \ , \tag{23}$$

for uMPS and uCPS, respectively, where n is the uCPS overlap. The entanglement entropy S is the simplest measure, for pure quantum states defined on a region $\mathcal{R}$, of entanglement between a sub-region $\mathcal{A} \in \mathcal{R}$ and the rest of the system. It is defined as:

$$S = -\mathrm{tr}(\rho_{\mathcal{A}} \log(\rho_{\mathcal{A}})) = -\sum_i \lambda_i^2 \log(\lambda_i^2) \ , \tag{24}$$

where λ_i are the Schmidt coefficients corresponding to the density matrix $\rho_{\mathcal{A}}$ associated with the sub-region $\mathcal{A}$. One can show that the Schmidt coefficients corresponding to a cut of the infinite spin chain into two semi infinite sub-chains are given by the singular values of $\rho_l^{\frac{1}{2}} \rho_r^{\frac{1}{2}}$, where ρ_l and ρ_r are the left and right uMPS environment matrices defined in Appendix A; for uCPS left and right environment eigenvectors see (5) and (6).

The convergence of S for a half-infinite chain, for the quantum Ising model in the paramagnetic phase (specifically for $J = 1$, $h = 0.5$), is depicted in the plot in Figure 7. We note that for the z-basis choice convergence of S approaches that of uMPS. This is not surprising, given that in the limit $h \to 0$ the ground state approaches a product state aligned along the z-basis, and so one would expect uCPS in this basis to be capable of capturing the exact state accurately already at small overlap.

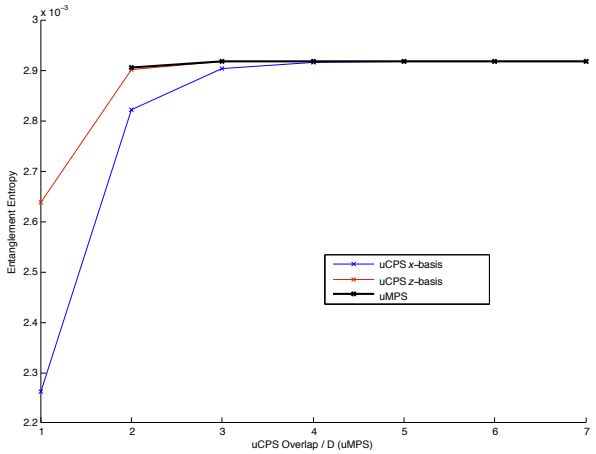

FIG. 7. (Colour online). Entanglement entropy S as function of uCPS overlap/uMPS bond dimension D, for the quantum Ising model in the ordered gapped phase ($J = 1$, $h = 0.5$).

The present example also demonstrates that uCPS is able to reproduce ground state energies to machine precision both for z- and x-basis choices. Due to the exponentially higher memory cost of uCPS compared to uMPS, in practice it is necessary to be far in the gapped phase to observe such convergence; for example, for $J = 1$, $h = 0.5$, machine precision is achieved for both basis choices with a 7-site uCPS overlap (which already corresponds to $D_{\mathrm{uCPS}} = 128$), while for $J = 1$, $h = 0.7$, the necessary overlap size is out of reach for a desktop with 16 GB of RAM. Remarkably, a naive first-order implementation of uCPS TDVP is sufficient to achieve this; despite the fact that most of the Gram matrix eigenvalues are zero to machine precision in this regime, no numerical instabilities are encountered. This behaviour depends crucially on implementing an appropriate pre-conditioner in the iterative subroutine step responsible for applying the inverse of the Gram matrix to a vector, as is described in Appendix C.

The convergence of the correlation length with uMPS bond dimension/uCPS overlap is depicted in the plot in Figure 8.

A notable feature is that, for both x- and z-basis choices, uCPS convergence is much smoother than for uMPS. Such behaviour is also observed at criticality, as is discussed later in this section (see Figure 10).

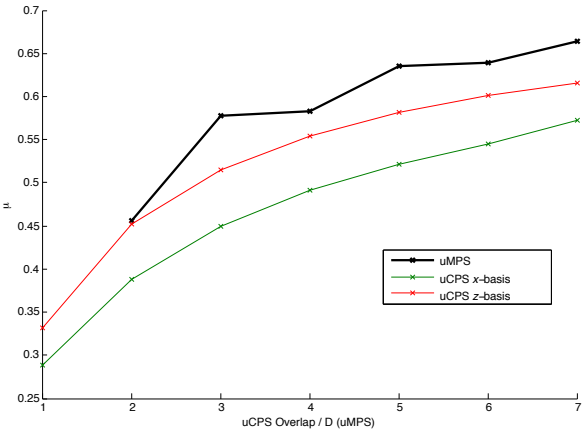

FIG. 8. (Colour online). Correlation length μ as function of uCPS overlap/uMPS bond dimension D, for the quantum Ising model in the ordered gapped phase ($J = 1$, $h = 0.5$).

Finally let us also note that tor the XY model, both at and away from criticality, convergence properties are similar to those of the quantum Ising model for γ close to one; as γ approaches zero convergence becomes slower, and the TDVP algorithm more sensitive to integration errors.

Finite entanglement scaling at criticality - Next we study the scaling properties of the correlation length (23) and entanglement entropy (24) for uCPS at criticality. Since the area law no longer holds, these quantities, while finite at any given bond dimension D, grow indefinitely as $D \to \infty$. Since a finite D bounds the amount of entanglement that can be encoded in a CPS/MPS, scaling with respect to D is referred to as *finite entanglement scaling*, and has been extensively studied in the uMPS setting [31, 34–36]. In what follows, we demonstrate that also at criticality overlap size plays the same role for uCPS as the bond dimension does for uMPS, and furthermore show how universal quantities can be calculated using uCPS finite entanglement scaling.

Let us first consider the scaling of the correlation length μ (23) with uCPS overlap. The work [31] provides numerical evidence that for a critical system uMPS scales with bond dimension as $(D_{\mathrm{uMPS}})^{\kappa}$ in the limit of large D_{uMPS}, where κ is a universal constant. In [34] it was furthermore argued that κ only depends upon the central charge c of the critical system via the relation:

$$\kappa = \frac{6}{c\left(\sqrt{\frac{12}{c}}+1\right)} \ . \tag{25}$$

The plot in Figure 9 demonstrates that, to high accuracy, for uCPS μ is instead proportional to $n^{\tilde{\kappa}}$, where n is the uCPS overlap, and $\tilde{\kappa}$ a constant. For the critical quantum Ising model, a linear fit of $\log(\mu)$ vs. $\log(n)$ for uCPS in the z-basis yields $\tilde{\kappa} = 1.004 \pm 0.006$. It seems plausible that the exact value of $\tilde{\kappa}$ is equal to one in the limit $n \to \infty$, i.e. that the correlation length is exactly proportional to uCPS overlap size. In the x-basis, as can be seen in Figure 9, a significant oscillation over the whole range of available n is present, and making any conclusion regarding the value of $\tilde{\kappa}$ in the limit $n \to \infty$ is difficult. Based upon the available data it is thus not possible to say whether or not $\tilde{\kappa}$ is basis dependent in this limit. This in turn makes it impossible to make any assertion as to whether the value of $\tilde{\kappa}$ is universal (and if so, in what sense).

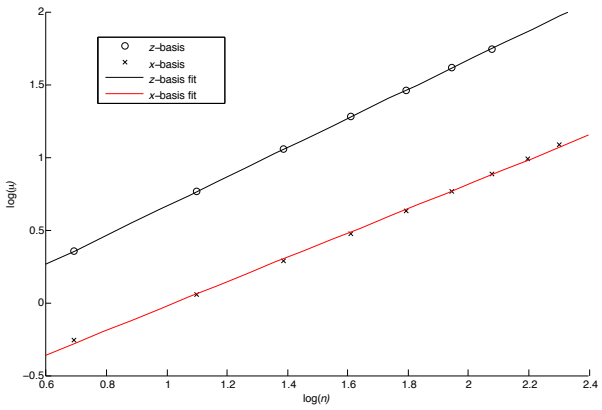

FIG. 9. (Colour online). The logarithm of the correlation length μ versus the logarithm of uCPS overlap for uCPS ground state approximations for the quantum Ising model at criticality, for both the x- and z- basis choices. In the z-basis the linear fit is particularly accurate, with slope equal to one to good approximation.

At present we can therefore not see any manner in which, in the uCPS context, something akin to (25) could be used to estimate the central charge c. Nevertheless, uCPS can be used to calculate universal quantities along the lines of [35, 36], where it has been demonstrated that a particularly powerful way to calculate critical exponents and the central charge is to scale not directly with respect to D, but with respect to the uMPS correlation length $\mu(D)$. The central charge, for example, can be estimated from the scaling of the entanglement entropy (27) with the correlation length (23) as follows. For a (1+1) critical system it has been shown [37, 38] that the entanglement entropy corresponding to an interval $\mathcal{A}$ of length $x_{\mathcal{A}}$ grows as:

$$S_{\mathcal{A}} = \frac{c}{3}\log(x_{\mathcal{A}}) + k \ , \tag{26}$$

where c is the central charge of the system, and k a constant. The entanglement entropy of the half-infinite line then scales as:

$$S = \frac{c}{6}\log(\mu) + \tilde{k} \ , \tag{27}$$

where μ is some length scale introduced in the system - in our case this is precisely the correlation length associated with finite bond dimension (23) - and $\tilde{k}$ is again some constant. Relation (27) in conjunction with (23) can thus be used to obtain an estimate of the central charge.

The scaling of S vs. $\log(\mu(n))$ for the quantum Ising model at criticality is depicted in the plot in Figure 10. As one can see, oscillations in the entanglement entropy as a function of $\log(\mu)$ observed for uMPS are absent from uCPS in a fixed basis. Performing a simple $1/n$, $n \to \infty$ extrapolation for the slope of this curve, with z-basis data up to $n = 10$, yields the estimate $c = 0.50142...$ for the central charge. This deviates from the exact value of $c = \frac{1}{2}$ only in the third digit,

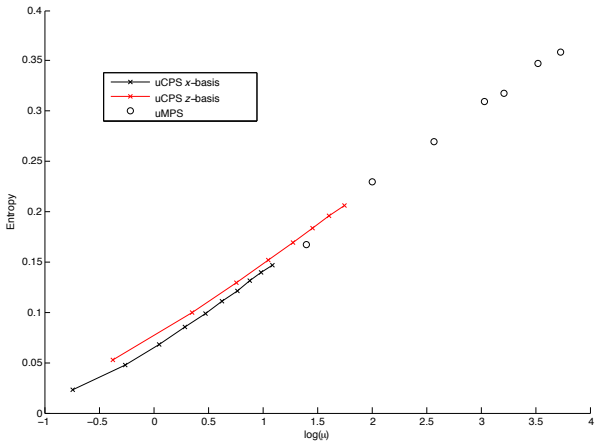

FIG. 10. (Colour online). The scaling of uCPS/uMPS entanglement entropy S vs $\log(\mu)$, where μ is the correlation length. The slope approaches a constant value, which is theoretically predicted to be $\frac{c}{6}$, where c is the central charge.

and provides better accuracy than the result obtained via finite uMPS entanglement scaling using all bond dimensions from $D = 2$ to $D = 64$ in [36]. The time needed to generate this uMPS data on an average spec current desktop, even when using more advanced techniques than the simple first-order implementation of imaginary time TDVP (such as the conjugated gradients method or iDMRG), is of the order of a week. On a comparably powerful computer the uCPS data used here was generated in a few hours. It should be noted, however, that uMPS finite entanglement methods are much more accurate for critical exponent than for the central charge estimates, and that the exponentially larger memory cost of uCPS means that uMPS is capable of accessing states with a lot more entanglement (as Figure 10 clearly demonstrates). Nevertheless, for the range of bond dimensions for which uCPS has a sufficient amount of RAM, the lack of oscillations means that uCPS can in practice be useful for making accurate estimates much more quickly than is possible with uMPS.

It is also interesting that uMPS and uCPS data points in Figure 10 lie roughly on the same line, and while uMPS data is significantly noisier, it seems to be bounded by the optimal z- and the suboptimal x- uCPS basis choices. This gives an indication that there may be some relationship between the oscillations in the uMPS data and the rotation invariance of the uMPS Ansatz.

In conclusion - While a single TDVP step scales better with bond dimension for uCPS than for uMPS, an exponentially larger amount of computer memory is needed in order to achieve the same accuracy with uCPS than with uMPS. However, the computational time required to obtain the same accuracy, surprisingly, scales in the same manner for both. The precise nature of the scalings is sensitively dependent upon both the model under investigation and the choice of basis. For models with degenerate ground states, convergence to the global minimum is in general only achieved for a certain fraction of TDVP runs. The uCSP algorithm exhibits certain advantages over uMPS when calculating universal quantities using finite entanglement scaling; oscillations found in uMPS are not present for uCPS, so scaling can be deduced more accurately with comparable computer time (but larger computer memory) cost. This seems to occur precisely *because* of the fixing of the basis. It would be particularly useful if oscillations are eliminated for CPS/string bond states in a similar manner beyond one dimension, where obtaining large number of points for a range of bond dimensions, as may be necessary with oscillations present, can be prohibitively expensive.

V. QUENCHES WITH uCPS

In this section we use real time TDVP applied to uCPS to study quenches across the critical point in the quantum Ising model (20) which exhibit so-called *dynamical phase transitions*. Dynamical phase transitions can occur whenever the return amplitude,

$$G(t) = \langle \Psi_0 | \, e^{-iHt} \, | \Psi_0 \rangle \ , \tag{28}$$

also referred to as the Loschmidt amplitude, has zeros [39]. The rate function for the return probability (referred to from here on just as the *rate function*):

$$l(t) = -\lim_{L\to\infty} \frac{1}{L} \log |G(t)|^2 \ , \tag{29}$$

can acquire non-analyticities, analogous to those present in the free energy in a thermodynamic setting. For a uCPS/uMPS, this quantity corresponds simply to the logarithm of the second largest eigenvalue of the transfer matrix.

The thermodynamic equivalent of (28) is obtained from purely imaginary time evolution, i.e. $t = -i\tau$ with τ real. In this case $G(\tau)$ corresponds to the canonical partition function of a system with finite length in the τ direction, and with boundaries described by $|\Psi_0\rangle$. *Equilibrium phase transitions* of the system are in correspondence with the zeros of $G(\tau)$ that occur as one takes the thermodynamic limit $L \to \infty$, which, if present, result non-analyticities in the free-energy $-\log|G(\tau)|^2$. It should be noted that, while the thermodynamic partition function can acquire zeros only in the thermodynamic limit, for the return amplitude (28) this can also happen also at finite system size. Thus, in general no simple correspondence between equilibrium and dynamical phase transitions exists [40].

The rate function (29) can be calculated exactly for the quantum Ising model, and a dynamical phase transition occurs for quenches across the critical point [41, 42], so in this setting a simple correspondence between the dynamical and equilibrium cases does exists. The ground state of the quantum Ising model (20) undergoes an equilibrium phase transition from a ferromagnetic phase for $h < 1$ to a paramagnetic phase for $h > 1$. This transition is captured accurately by uMPS even at small bond dimension, with the expectation value of order parameter operator $< \hat{\sigma}^z >$ going from a positive value for $h < h_c$, to zero[12] for $h > h_c$, with h_c approaching the exact value $h_c = 1$ from above with increasing bond dimension. The uCPS behaviour is highly basis-dependent, as is demonstrated by the plot in Figure 11. The critical point is approximated with comparable accuracy by uMPS and uCPS in the z-basis, at bond dimension equal to uCPS overlap, but

[12] To machine precision.

the uCPS approximation gets increasingly less accurate for choices of basis away from z, and completely misses the phase transition in the x-basis.

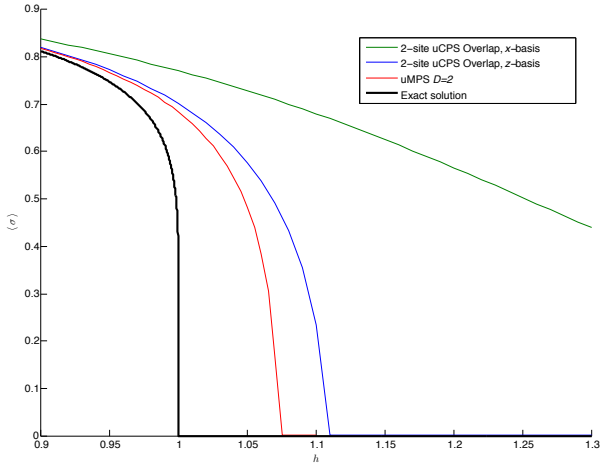

FIG. 11. (Colour online). Both uCPS and uMPS are capable of capturing the quantum Ising model equilibrium phase transition, from the ferromagnetic phase, $h < 1$, characterised by a non-zero value for the order parameter operator expectation value $< \sigma^z >$, to the paramagnetic phase $h > 1$ when $< \sigma^z >= 0$. The accuracy of the uCPS approximation for the critical point, $h = 1$, degrades as a function of the rotation angle away from the z-basis. In the x-basis the phase transition is entirely missed.

In the following, we concentrate on a quench initiated in the paramagnetic phase, with $h = 1.5$ in the quantum Ising Hamiltonian (20), with the time evolution performed according to a Hamiltonian with the magnetic field (instantaneously) changed to $h = 0.1$; the post-quench Hamiltonian is therefore deep in the ferromagnetic phase. The reversed quench, from the ferromagnetic to the paramagnetic phase, also exhibits a dynamical phase transition, but is more unwieldy to analyse [42]. As it yields very similar conclusions regarding the properties of uCPS, it will not be explicitly discussed here.

In the z-basis uCPS captures the dynamical phase transition very accurately, as demonstrated in the plot in Figure 12. In addition to capturing non-analyticities, simulations with uCPS in this basis exhibit approximate recurrences at large times, i.e. beyond the point at which the exact rate function is captured correctly, for all values of uCPS overlap. For single-site overlap recurrences are in fact exact. Plots of uCPS approximations for single- and 3-site overlap in the z-basis are depicted in the two top plots of Figure 14 and clearly show recurrences.

In contrast, uCPS in the x-basis (Figure 13) do not capture the non-analyticities in the rate function at all. For times prior to the first non-analyticity, the accuracy of the uCPS approximation of the rate function increases with increasing overlap, however the non-analyticity is never actually captured. Beyond the point at which the non-analyticity occurs in the exact function, the behaviour of the uCPS approximation completely misses the correct behaviour and is chaotic, not converging to any definite function with increasing uCPS overlap size. It should be noted that non-analytic behaviour is not observed at any time in the x-basis uCPS approximation. This is very different to the large time behaviour of uCPS in the z-basis.

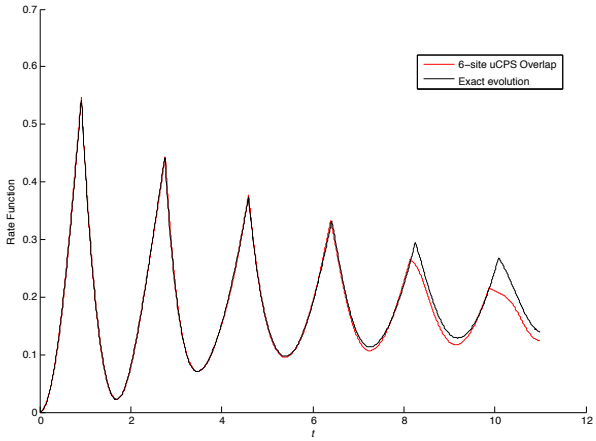

FIG. 12. (Colour online). The rate function vs. time in the z-basis for 6-site uCPS overlap compared with the exact evolution. The quench corresponds to the ground state of the quantum Ising Hamiltonian with $h = 1.5$ (paramagnetic phase) evolved with the $h = 0.1$ Hamiltonian (ferromagnetic phase).

The behaviour of the uMPS approximation of the quench, at large times, exhibits a combination of that observed for z- and x-basis uCPS approximations in the following sense: as can be seen in lowermost plot in Figure 14, beyond the point at which the exact rate function is captured accurately (for a given bond order), the uMPS behaviour is chaotic, as is the case for uCPS in the x-basis, but it does exhibit non-analyticities at arbitrarily large times, a feature observed for uCPS in the z-basis.

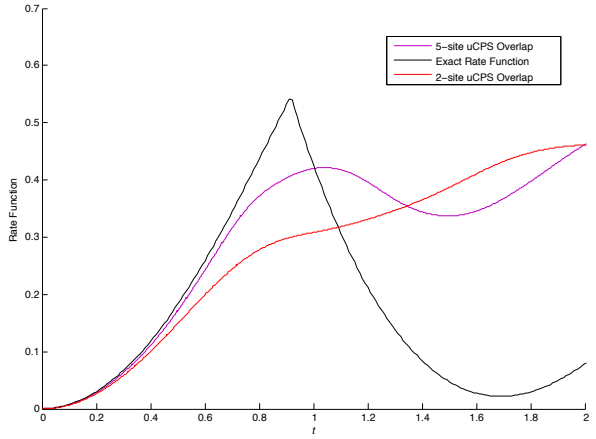

FIG. 13. (Colour online). The uCPS approximation, for 2- and 5-site uCPS overlap in the x-basis, of the rate function for the $h = 1.5 \rightarrow h = 0.1$ quench in the quantum Ising model. In this basis uCPS completely misses the dynamical phase transition. The behaviour of the uCPS approximation is analytical even at large times (not displayed).

Comparing the $D = 6$ uMPS approximation in Figure 14 with the 6-site uCPS overlap approximation in the z-basis in Figure 12 demonstrates that the quality of the uCPS and uMPS approximations is comparable. This is another example - here in the context of real time evolution - of the general observation that the capacity of uCPS and uMPS to cap-

ture the properties of spin systems is roughly the same, in an optimal basis, for uCPS overlap equal to the uMPS bond dimension. In addition, we observe that, for the examples studied in this section, the time needed to run uMPS and uCPS approximations of the same quench at uCPS overlap size equal to uMPS bond dimension is roughly the same. This demonstrates that also in the context of quenches it is only the computer memory cost, and not the computer time cost, that scales exponentially worse for uCPS.

VI. DISCUSSION

This paper has centred around developing the deterministic TDVP algorithm to study a restricted class of uCPS states. Such an algorithm is itself a departure as CPS states are usually optimised stochastically and - as far as we are aware - not used to study real time Hamiltonian evolution. Our analysis has revealed several interesting features of the numerical cost of using uCPS and their sensitivity to choice of basis.

Computational cost of uCPS TDVP algorithm - implementation of the TDVP algorithm shows interesting efficiency gains, aspects of which may be extendible to higher dimensions. In Section IV it has been demonstrated that the cost of a single uCPS TDVP step scales, optimally, as $O(D^3)$. Since uCPS maps to an uMPS with bond dimension equal to the physical spin dimension, and the cost of one uMPS TDVP step is $O(dD^3)$, one naively expect the cost of each uCPS TDVP step to scale as $O(D^4)$, but the special structure of uCPS enables this to be improved upon. This special structure is most effectively analysed by considering the uCPS $\rightarrow$ uMPS mapping, which necessitates the introduction of the so called copier tensor, and the fact that a singe n-site coper factorises into a product of n single-site copiers already implies the reduction to $O(\log(D)D^3)$. The mapping of a generic CPS to PEPS involves copiers that factorise in a similar manner, so it is likely that an increase in efficiency is possible also in more than one dimensions for deterministic algorithms that respect the CPS structure.

However, the efficiency with which a uCPS represents a state to a given degree of accuracy is not as great as one might naively anticipate. As discussed in Section II, we have considered a restricted set of CPS that are equivalent to MPS with bond and local Hilbert space dimension equal. One might expect then that uCPS with an optimally chosen basis are of comparable accuracy to uMPS of this bond order. This is not the case. Both deep in the gapped phase and also at criticality, quantities of interest, such as the ground state energy, converge with size of the uCPS overlap n in roughly the same manner as they converge with bond dimensions for uMPS. The implication is that, as far as physical observables are concerned, uCPS scaling is exponentially worse than uMPS, since $D_{\text{uCPS}} = s^n$, where s is the dimension of the spin of a single site. For critical systems one can make a particularly precise statement: while for uMPS the correlation length μ increases with bond dimension D as $\mu \propto D^{\kappa}$, in the limit $D \rightarrow \infty$, for uCPS the correlation length scales instead as $\mu \propto n^{\tilde{\kappa}}$, where κ and $\tilde{\kappa}$ are constants. It is not possible to determine, from the accessible range of uCPS overlap sizes, whether or not $\tilde{\kappa}$ is basis dependent, and thus, whether or not it encodes universal behaviour. Interestingly for the quantum Ising model in the z-basis $\tilde{\kappa}$ seems to be one to good accuracy, and so the correlation length is precisely proportional to the uCPS overlap.

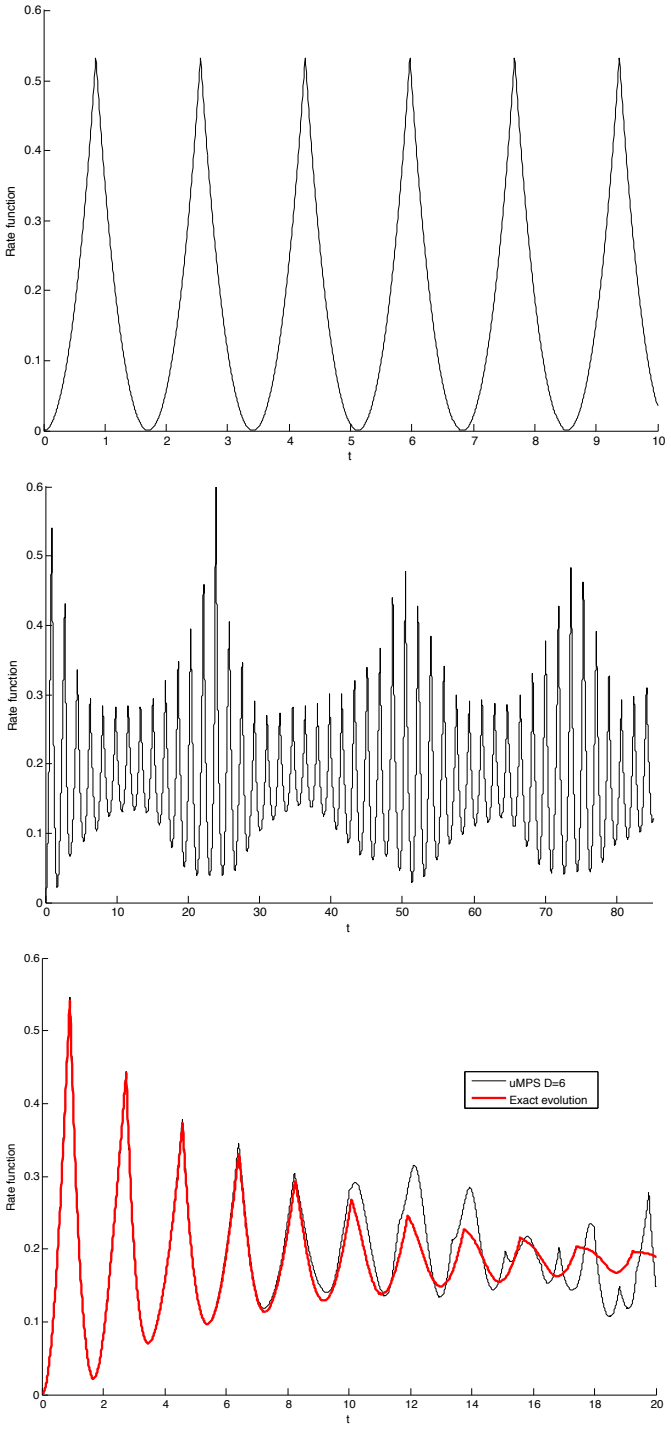

FIG. 14. (Colour online). Rate function vs. time in z-basis exhibits recurrence when projected to the uCPS manifold. Recurrence is exact for 1-site CPS overlap (top), and approximate when the overlap is larger than one, as demonstrated for the 3-site uCPS overlap case (middle). The uMPS approximation does not exhibit recurrences at large times; beyond the point at which the rate function is accurately captured, the behaviour is chaotic (bottom), and the precise behaviour is also highly bond-dimension dependent.

Surprisingly, while the *memory* cost is exponentially worse

for uCPS than for uMPS, the computer *time* needed to fully converge to the optimal uCPS ground state approximation, using imaginary time TDVP initiated from a random state, scales in the same way for uCPS with respect to overlap, as it does for uMPS with respect to bond dimension. Combined with the above result, that uCPS achieves roughly the same accuracy as uMPS for overlap equal to uMPS bond dimension, this observation implies that uCPS in the optimal basis matches uMPS accuracy with the same computer time cost, but an exponentially worse memory cost. This does not contradict the statement that the cost of a single step in the imaginary time TDVP is exponentially worse for uCPS, both in time and memory cost (since at each step exponentially larger matrices need to be multiplied). The speedup observed with uCPS, over the course of the whole imaginary time TDVP run, reflects the fact that for uCPS much larger time steps can be taken than for uMPS before instability sets in. In addition, while for imaginary time TDVP integration errors due to excessively large time steps can conspire to actually aid convergence, both for uCPS and uMPS - and of course as long as these time steps are not too large - this effect aids convergence more efficiently in the case of uCPS. The net result is that the the exponentially higher cost of a single step for uCPS in the TDVP algorithm is counteracted by these two effects to bring the computer time cost of the whole TDVP run down, so that it actually scales exponentially better than would be expected when considering the costs of a single step of the algorithm; clearly, none of this improves the memory requirements.

The computational costs are observed to have the same properties in the context of real-time TDVP, which has been used to simulate quantum quenches in Section V. The only caveat is that in the analysis, in place of considering the time necessary for imaginary time TDVP to converge, one must compare the uCPS vs. uMPS cost of simulating the evolution of a quench over some reasonably long time. Given that integration errors are not of great concern for imaginary time TDVP, only a simple Euler-step integrator was used to generate the ground state approximations, and one may object that some of the conclusions made above regarding the cost of imaginary time TDVP are merely artefacts of the particularly simple type of integration scheme. This is not the case, as can be checked by employing a more sophisticated integration algorithm. Moreover, our real time uCPS quench simulations were computed using the adaptive step Runge-Kutta-Fehlberg 4(5) method, with the same conclusions regarding uCPS vs. uMPS computational costs.

Basis dependence - In addition to the above, basis dependence has a major impact on the behaviour of the uCPS TDVP algorithm. In Section V it was demonstrated that with a sub-optimal choice of basis important physics can be missed, as demonstrated by the fact that both dynamical and static phase transitions in the quantum Ising model are not observed if one chooses to work in the x-basis. Clearly, this is a cautionary lesson when working with higher dimensional CPS/string bond states.

In the imaginary time context, it is observed that the uCPS TDVP algorithm converges to the global minimum irrespective of the choice of the initial random state in general only when the ground state is not degenerate. For degenerate vacua the algorithm converges to the global minimum only for a certain fraction of runs initiated from a random state, except possibly for very special choices of basis for which convergence is unique. If there is a finite number of degenerate vacua, the number of possible minima is in general equal to the degeneracy and does not increase with bond dimension. For a continuous symmetry group uCPS TDVP seems to converge - as far as our numerical analysis is capable of ascertaining - to a number of local minima that increases without bound with uCPS overlap. The statements concerning the scaling and cost of obtaining a ground state made above still hold, but when convergence is not unique the algorithm acquires a probabilistic ingredient, and for continuous symmetries the scaling of the computational cost can increase by more than a constant factor.

In reference to models examined in Section IV, for the quantum Ising Hamiltonian entanglement is generated by the single term $\hat{\sigma}^z \otimes \hat{\sigma}^z$, and a number of suitable basis choices exist that yield unique convergence; for the XY model with $\gamma = 0$, which has $U(1)$ symmetry, no such choice exists. At the other extreme convergence is observed to be unique, irrespective of basis choice, for the Heisenberg model, which is described by a rotation invariant Hamiltonian. Here all basis choices are clearly equivalent, and imaginary time TDVP always converges to a unique minimum. However, energy converges extremely slowly compared to the other models studied (see Figure 6) which demonstrates that uCPS is best suited to the study of Hamiltonians maximally aligned with the uCPS basis.

In conclusion, a good basis choice is one for which the entanglement generating terms are 'optimally aligned', in some sense, with the uCPS basis. The problem of what precisely is meant by 'optimal' is encountered for all CPS/string bond state approaches, and is a difficult one to tackle with any generality - not only due to technical challenges, but also because the answer depends upon what precisely one wishes to achieve. We have illustrated how this pans out in detail in the context of uCPS in Section IV. In this case, convergence properties, probabilistic aspects of the algorithm, and the accuracy of estimates for physical observables are all basis dependent, yet can not in general be optimised simultaneously; the choice depends upon which of these properties one wishes to prioritise.

VII. CONCLUSIONS

Although quantum states in dimensions higher than one can be represented efficiently by tensor networks, physical properties may not in general be calculated efficiently without further approximations. One way around this is to place additional restrictions upon the tensor network so that its properties are easier to calculate. Such restrictions inevitably involve compromises and a balance between efficiency gains and accuracy. We have investigated this balance in the controlled context of a restricted class of uniform one-dimensional Correlator Product States that may also be considered a restriction upon uniform matrix product states. Similar restrictions may be applied in higher dimensions - Correlator Product States with small, double overlaps can be mapped to small bond order PEPS.

Our main results are:

- The application of the time dependent variational principle to uCPS. Usually CPS - as well as the more general class of string bond states - are optimised using a stochastic Monte-Carlo type approach. They are well-suited to this because of their efficient sampleability.

By considering the mapping of uCPS into uMPS, one naively expects a single uCPS TDVP step to scale as $O(D^4)$. Utilising the special structure uCPS, we have shown that the cost of a single uCPS TDVP step can be reduced to $O(D^3)$. Since this is based upon the properties of the copier tensor whose properties generalise to higher dimensional CPS $\rightarrow$ PEPS mappings, our analysis indicates that a similar reduction should be possible for higher dimensional deterministic algorithms that respect the CPS structure.

- The capacity of uCPS to capture physical information about a system scales exponentially worse with bond dimension than uMPS. In order to capture the ground state energy, or a quantum quench, to the same accuracy as a D dimensional uMPS, one has to work with uCPS with overlap size of the order $n \approx D$, which is exponentially more expensive since $D_{\mathrm{uCPS}} = s^n$, where s is the dimension of a single spin. This is, surprisingly, only reflected in computer memory usage, not in the computer time needed to obtain the ground state via imaginary time TDVP, or to run quantum quenches, which for a uCPS with overlap of size n is of the same order as for uMPS with $D = n$.
- The choice of uCPS basis has a strong effect upon the behaviour of the TDVP algorithm, on the capacity of uCPS to accurately approximate the physical system under consideration, as well as on the behaviour of uCPS under scaling. A good basis choice, generally speaking, has the property that it is closely aligned with the entanglement generating terms in the Hamiltonian. With an optimal choice of basis uCPS will generally capture the physics as well as uMPS, for uCPS overlap equal to the uMPS bond dimension. On the other hand, with a suboptimal choice uCPS can completely fail to capture important physics such as equilibrium or dynamical phase transitions. We also observe that, being an Ansatz that is not rotation invariant, except for special basis choices uCPS will in general break any degeneracy present in the exact ground state (or in the related uMPS approximation): depending on details of the random initial state, imaginary time TDVP will converge to local minima associated with this separation of otherwise degenerate energy levels.
- Having fixed a basis, the scaling of a physical quantity with bond dimension is much smoother for uCPS than for uMPS. In particular, the scaling of quantities such as entanglement entropy or the correlation length often has strong oscillations at lower bond dimensions in the case of uMPS, and these almost entirely disappear for uCPS in a fixed basis.
- Some properties of uMPS exhibit a combination of features that can be isolated by making judicious choices of uCPS basis. For example, the oscillations seen in the scaling behaviour of uMPS seem to be bounded by the smooth behaviour of uCPS scalings, at one end by the uCPS in the optimal basis and the other by the least optimal choice of basis (see Figure 10). Similarly the behaviour of uMPS at large times for the quench exhibiting dynamical phase transitions is a combination of the recurring non-analyticities seen for the optimal uCPS choice of basis, and the chaotic but analytic behaviour observed in the least optimal basis (see Figures 13 and 14).

The analysis in this paper has mostly been geared towards bettering our theoretical understanding uCPS compared to standard uMPS, with a view to identifying characteristic properties of the uCPS variational manifold that may be of use when studying CPS/string bond states in general, and in particular in higher dimensions. It should be stressed that uCPS has potential practical advantages already in one dimension. A general observation made at various points in this paper is that the uCPS TDVP algorithm, applied to a suitable Hamiltonian and in an optimal basis, is very robust - both in its imaginary and real time variants (see e.g. Figures 7 and 14) - under the right circumstances more so than a comparably costly uMPS TDVP run. A further advantage of uCPS is described at the end of Section IV, where it is shown that one can utilise the superior scaling properties of uCPS compared to uMPS in order to generate estimates of universal quantities in critical theories, with accuracies not achievable with comparable computational time cost using uMPS. Finally, an aspect of our analysis that has not been emphasised in the course of this paper is that, while uCPS yields similarly accurate estimates of physical quantities for overlap sizes equal to uMPS bond dimension, the actual bond dimensions accessed by uCPS are exponentially larger than what is accessible with uMPS at comparable computational time cost. For example, for the critical quantum Ising model the bond dimension $2^{10} = 1024$ uCPS ground state estimate in the x-basis is reached with roughly the same computer time cost needed to generate the uMPS $D = 10$ ground state approximation (see Figure 4). One interesting question is, for example, whether the uCPS $\rightarrow$ uMPS mapping described in Section II could provide a more efficient way of initialising a $D = 1024$ iDMRG run, than e.g. by building it up from a $D = 10$ uMPS state?

ACKNOWLEDGMENTS

We gratefully acknowledge support by EPSRC under grant number EP/I031014/1 and EP/I004831/2. T. Đurić acknowledges support from the EU Grant QUIC (H2020-FETPROACT-2014, Grant No. 641122). V. S. would like to thank Damian Draxler, Jutho Haegeman, Michaël Mariën, and Frank Verstraete for helpful discussions.

Appendix A: Review of Matrix Product States and the Time-Dependent Variational Principle

A matrix product state (MPS) takes the form:

$$|\Psi[A]\rangle = \sum_{i_1,i_2,\cdots,i_N}^{d} v_L^\dagger A_1^{i_1} A_2^{i_2} \cdots A_N^{i_N} v_R \, |i_1, i_2, \cdots, i_N\rangle \ , \tag{A1}$$

where d is the number of physical (spin) degrees of freedom, and A^{i_a} is a $D_{A-1} \times D_A$ matrix, while v_L and v_R are vectors of dimensionality D_0 and D_N respectively. The dimension of these internal/virtual indices is referred to as the *bond dimension*. The open boundary condition Ansatz is presented here for the sake of concreteness - the periodic boundary condition MPS Ansatz corresponds to taking a trace in place of contraction by $v_L^\dagger$ and v_R. Graphically the state (A1) can be

represented as:

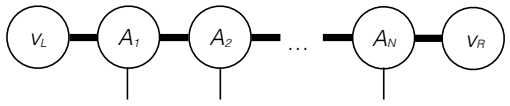

,

where the thick lines represent the internal bond dimension indices, while the physical spins are denoted by the thin lines with 'uncontracted' ends. Here and in what follows we will not graphically depict any variation in the dimensionality of physical or bond indices.

Transformations of the MPS matrices of the form:

$$A^{i_n} \to G_{n-1} A^{i_n} G_n^{-1} \ , \tag{A2}$$

where $G_n, G_{n-1} \in GL(D, \mathbb{C})$, leave the state invariant, and are referred to as gauge transformations. The gauge transformations of the boundary (co)-vectors are given by $v_L^\dagger \to v_L^\dagger G_0^{-1}$, and $v_R \to G_N v_R$.

Applying the time-dependent variational principle (TDVP; see introduction to Section III) to the MPS variational class [28, 30] yields:

$$\dot{A}^*(t) = \underset{\dot{A}(t)}{\arg\min} \left\| \left| \Phi(\dot{A}, A) \right\rangle + i\hat{H}(t) \left| \Psi(A(t)) \right\rangle \right\| \ , \tag{A3}$$

where the object $|\Phi(dA, A)\rangle$ is given by:

$$\begin{aligned} |\Phi(dA, A)\rangle = & \sum_{i_1, \cdots, i_N} v_L^\dagger dA_1^{i_1} A_2^{i_2} \cdots A_N^{i_N} v_R \left| i_1, i_2, \cdots, i_N \right\rangle \\ & + \sum_{i_1, i_2, \cdots, i_N}^{d} v_L^\dagger A_1^{i_1} dA_2^{i_2} \cdots A_N^{i_N} v_R \left| i_1, i_2, \cdots, i_N \right\rangle \\ & + \quad \cdots\cdots \\ & + \sum_{i_1, i_2, \cdots, i_N}^{d} v_L^\dagger A_1^{i_1} A_2^{i_2} \cdots dA_N^{i_N} v_R \left| i_1, i_2, \cdots, i_N \right\rangle \\ & := \sum_\alpha dA^\alpha \frac{\partial}{\partial A^\alpha} \left| \Psi(A) \right\rangle \equiv \sum_\alpha dA^\alpha \left| \partial_\alpha \Psi(A) \right\rangle \ . \end{aligned} \tag{A4}$$

The α index on the last line combines physical, virtual, and site indices. Unlike elements in the variational class of MPS, the objects $|\Phi(dA, A)\rangle$ form a vector space - and are in fact tangent vectors to the MPS manifold [28, 43]. The gauge invariance of MPS (A2) can be shown to imply the invariance of tangent states under

$$dA^{i_n} \to dA^{i_n} + X_{n-1} A^{i_n} - A^{i_n} X_n \ , \tag{A5}$$

where the matrices X live in the Lie algebra of the gauge group.

In this paper we will study MPS directly in the thermodynamic limit ($N \to \infty$), and will use a translation invariant Ansatz, taking the MPS tensors to be position independent (which clearly requires a constant bond dimension D). We refer to the class of such states as uniform MPS (uMPS). Graphically a uMPS is represented as:

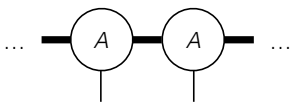

.

The largest eigenvalue of the uMPS transfer matrix E,

$$E := \sum_i A^i \otimes \overline{A}^i \ , \tag{A6}$$

needs to be fixed to unity in order to ensure finite normalisation. The state norm is given by $(l|r)$, where $(l|$ and $|r)$ are respectively the left and right eigenvectors of E corresponding to eigenvalue 1, and since the eigenvectors can be rescaled freely the state can always be normalised to one.

The expectation value of a local operator acting on n sites can now be written as:

$$\begin{aligned} & \langle \Psi[A] | \hat{O} | \Psi[A] \rangle = \\ & \sum_{\substack{i_1, i_2, \cdots, i_n, \\ j_1, j_2, \cdots, j_n}}^{d} (l| O_{i_1 \cdots i_n}^{j_1 \cdots j_n} \left(A^{i_1} \cdots A^{i_n} \right) \otimes \left(\overline{A}_{j_1} \cdots \overline{A}_{j_n} \right) |r) \ . \end{aligned} \tag{A7}$$

The cost of computing this object, using the optimal sequence of contractions, can be seen to scale as $O(ndD^3)$.

The uMPS parametrisation is invariant under gauge transformations:

$$A^i \to G A^i G^{-1} \ , \tag{A8}$$

and these can be used to set either $\rho_l = \mathbb{1}$ or $\rho_r = \mathbb{1}$ (but in general not both), where ρ_l is the D^2 dimensional co-vector $(l|$ reshaped to a $D \times D$ matrix, and similarly for ρ_r. The corresponding gauge transformations in the tangent plane (see (A5)) are given by $dA^i \to dA^i + XA^i - A^i X$, and this gauge freedom can be used to set either:

$$(l| \sum_i dA^i \otimes \overline{A}^i = 0 \quad \text{or} \quad \sum_i dA^i \otimes \overline{A}^i |r) = 0 \ , \tag{A9}$$

but again in general not both. These are referred to as, respectively, the left and right tangent space gauge conditions.

The Gram matrix, given by:

$$G_{\overline{\alpha}\beta} := \langle \partial_{\overline{\alpha}} \Psi(A) | \ \partial_\beta \Psi(A) \rangle \ , \tag{A10}$$

defines a natural metric on the uMPS manifold [43]. The TDVP equations can be formally expressed as:

$$\sum_\beta G_{\overline{\alpha}\beta} \dot{A}^\beta = -i \langle \partial_{\overline{\alpha}} \Psi(A) | \hat{H}(t) | \Psi(A(t)) \rangle \ . \tag{A11}$$

Imposing either the left or right tangent gauge condition (A9) simplifies the equations significantly, and is in fact necessary to eliminate infinities in (A11) stemming from transformations along the uMPS state itself. The expression for the overlap of two tangent vectors takes the simple form:

$$\sum_{\overline{\alpha}, \beta} \overline{A}'^{\overline{\alpha}} G_{\overline{\alpha}\beta} dA^\beta = |\mathbb{Z}| (l| dA \otimes d\overline{A}' |r) \ , \tag{A12}$$

since the left (or right) gauge condition implies that all terms for which the dA and $d\overline{A}'$ tensors are not at the same site are zero. The Gram matrix then takes the simple form:

$$G = \rho_l \otimes \rho_r \ . \tag{A13}$$

The nature of the right hand side of (A11) is elucidated by contracting it with $d\overline{A}'^{\alpha}$:

$$\begin{aligned} & \sum_{\overline{\alpha}} d\overline{A}'^{\overline{\alpha}} \langle \partial_{\overline{\alpha}} \Psi(A) | \hat{H}(t) | \Psi(A(t)) \rangle = \\ & |\mathbb{Z}| \Big((l| H^{AA}_{d\overline{A}'\overline{A}} |r) + (l| H^{AA}_{\overline{A} d\overline{A}'} |r) \\ & + (l| H^{AA}_{\overline{A}\overline{A}} (\mathbb{1} - E)^{PI} (A \otimes d\overline{A}') \| r) \Big) \ . \end{aligned} \tag{A14}$$

Here $H^{AA}_{\overline{A}d\overline{A}'}$, for example, stands for the contraction of the two 'ket' indices of a local term in the Hamiltonian, which is assumed to be translation invariant, with two A uMPS tensors, and the contraction of its 'bra' indices with $\overline{A}$ and $d\overline{A}'$ tensors. PI indicates a pseudo-inverse on the subspace of $D^2 \times D^2$ matrices defined by the projector $\mathbb{1}_{D^2\times D^2} - |r)(l|$. This term stems form a summation over all contributions with the $d\overline{A}'$ tensor to the right of the local Hamiltonian terms (the terms when $d\overline{A}'$ is on the left are zero due to the left tangent gauge condition, which is assumed here).

Appendix B: Contracting uCPS with computational cost $O(D^3)$

Let us first consider the following term appearing on the right hand side of (16) (i.e. in $-i\left\langle \partial_{\overline{\tilde{\alpha}\beta}}\Psi(C(t))\right| \hat{H}(t) \left|\Psi(C(t))\right\rangle$):

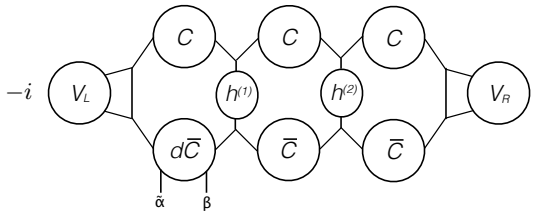

where a two-site Hamiltonian has been assumed. Symbolically this corresponds to:

$$\begin{aligned} & (C_{jm}) \\ T^{\tilde{\alpha}\beta} := -iV_{(L)i}C_{ij}\left(h^{(1)}_{jk}h^{(2)}_{mn}\right)C_{ma}V_{(R)a} \quad , \\ & d\overline{C}^{\tilde{\alpha}\beta}_{ik}\left(\overline{C}_{kn}\right)\overline{C}_{na} \end{aligned} \tag{B1}$$

where a summation over all repeated indices is implied, $h^{(1)}$ and $h^{(2)}$ constitute a contribution to a two-site term of the Hamiltonian, and

$$d\overline{C}^{\tilde{\alpha}\beta}_{ik} = V^{\tilde{\alpha}}_{(R)i}V^{\beta}_{(L)k}(.*)\left(1(./)\overline{C}_{ik}\right) \ . \tag{B2}$$

The following contraction ordering (which starts from the right):

1. $T^{(1)}_{mn} := C_{ma}\overline{C}_{na}V_{(R)a}$ is obtained at cost $O(D^3)$,
2. $T^{(2)}_{mn} := h^{(2)}_{mn}(.*)T^{(1)}_{mn}$ is obtained at cost $O(D^2)$,
3. $T^{(3)}_{jk} := C_{jm}\overline{C}_{kn}T^{(2)}_{mn}$ is obtained at cost $O(D^3)$,
4. $T^{(4)}_{jk} := h^{(1)}_{jk}(.*)T^{(3)}_{jk}$ is obtained at cost $O(D^2)$,
5. $T^{(5)}_{ik} := [V_{(L)}(.*)C]_{ij}T^{(4)}_{jk}$ is obtained at cost $O(D^3)$,
6. $T^{(6)}_{ij} := T^{(5)}_{ik}(.*)\left(1(./)\overline{C}_{ik}\right)$ is obtained at cost $O(D^2)$,
7. $T^{\tilde{\alpha}\beta} = -iT^{(6)}_{ik}V^{\tilde{\alpha}}_{(R)i}V^{\beta}_{(L)k}$ is obtained at cost $O(D^3)$,

can be seen to yield maximal cost $O(D^3)$ at any internal step. It is easy to find other contraction orderings that yield the same cost.

There are two more contributions to the right hand side of (16) that are similar to the above. One of these is obtained by substituting $d\overline{C}^{\tilde{\alpha}\beta}_{kn}$ in place of $\overline{C}_{kn}$ in (B1), and the other by making the same substitution in place of $\overline{C}_{na}$. A contraction ordering with $O(D^3)$ efficiency can be obtained for these in a very similar manner as what has been demonstrated above.

The remaining contribution to (16) corresponds to a sum over all terms with $d\overline{C}^{\tilde{\alpha}\beta}$ not coinciding with the Hamiltonian, and is given by:

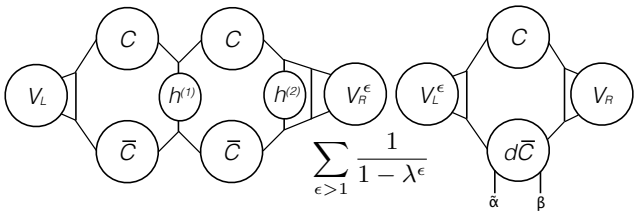

The uMPS expression corresponding to this is given by the term on the last line of (A14); as with uMPS, the terms with $d\overline{C}$ to the left of the Hamiltonian are zero due to the left tangent gauge condition, which has been assumed here. It is worth noting that while the computation of the pseudo-inverse acting on a vector in the uMPS expression can only be achieved at cost $O(dD^3)$, by recourse to an iterative procedure such as the biconjugate gradient algorithm, the uCPS term above can be computed explicitly at cost $O(D^3)$ using a contraction ordering similar to that described for (B1).

Appendix C: uCPS Gram matrix inverse pre-conditioning

The uCPS Gram matrix is given by:

$$G^{\tilde{\gamma}\delta|\tilde{\alpha}\beta} = \sum_{ij}\lambda_{ij}T^{\tilde{\gamma}\delta|\tilde{\alpha}\beta}_{ij} \ , \tag{C1}$$

where

$$T^{\tilde{\gamma}\delta|\tilde{\alpha}\beta}_{ij} = \left(\overline{V}^{\tilde{\gamma}}_{(R)i}\overline{V}^{\delta}_{(L)j}\right)(.*)\left(V^{\tilde{\alpha}}_{(R)i}V^{\beta}_{(L)j}\right) \ . \tag{C2}$$

The convention is that the non-tilde indices run over the whole range, 1 to D, while the tilde-indices correspond to a truncation and run from 2 to D, and $V^1_{(L)} \equiv V_{(L)}$ and $V^1_{(R)} \equiv V_{(R)}$ are the left and right eigenvectors corresponding to eigenvalue one.

Let us first consider the matrix:

$$\tilde{G}^{\tilde{\gamma}\delta|\tilde{\alpha}\beta} = \sum_{ij}(1(./)\lambda_{ij})\tilde{T}^{\tilde{\gamma}\delta|\tilde{\alpha}\beta}_{ij} \ , \tag{C3}$$

where

$$\tilde{T}^{\tilde{\gamma}\delta|\tilde{\alpha}\beta}_{ij} = \left(\overline{V}^{\tilde{\gamma}}_{(L)i}\overline{V}^{\delta}_{(R)j}\right)(.*)\left(V^{\tilde{\alpha}}_{(L)i}V^{\beta}_{(R)j}\right) \ . \tag{C4}$$

The obstruction to $\tilde{G}$ being the inverse of G can be understood as originating in the truncation of the eigenvalue one eigenvector in our implementation of the left tangent gauge condition (14). Namely, the eigenvector matrices obey:

$$\begin{aligned} V^{\delta}_{(L)i}V^{\epsilon}_{(R)i} &= \delta^{\delta\epsilon} \ , \\ \overline{V}^{\epsilon}_{(R)j}V^{\epsilon}_{(L)m} &= \delta_{jm} \ , \\ V^{\tilde{\delta}}_{(R)i}V^{\tilde{\epsilon}}_{(L)i} &= \delta^{\tilde{\delta}\tilde{\epsilon}} \ , \\ \overline{V}^{\epsilon}_{(L)_j}V^{\epsilon}_{(R)m} &= \delta_{jm} \ , \\ V^{\tilde{\gamma}}_{(L)j}\overline{V}^{\tilde{\gamma}}_{(R)m} &= \delta_{jm} - V_{(L)j}V_{(R)m} \ , \\ \overline{V}^{\tilde{\gamma}}_{(R)j}V^{\tilde{\gamma}}_{(L)m} &= \delta_{jm} - V_{(R)j}\overline{V}_{(L)m} \ , \end{aligned} \tag{C5}$$

where summation over repeated indices is understood. The presence of projectors in the last two lines expresses the deformation of the exact orthogonality relations due to the $\alpha \to \tilde{\alpha}$ truncation, which prevents (C3) from being the Gram matrix inverse.

In order to explicitly compute the action of the inverse of G on a vector with computational cost scaling as $O(D^3)$, it is necessary to be able to express G^{-1} in the general form (C3), or at least as a sum of a constant number of terms of this form. We have not managed to find any suitable solution that would bypass the obstruction described above (but have also not proved that doing so is impossible). However, in order to achieve $O(D^3)$ scaling, one can instead make use of an iterative algorithm, the biconjugate gradient (stabilised) method. This algorithm provides a solution for $\vec{x}$ in the equation $\mathbb{A}\vec{x} = \vec{b}$, where $\mathbb{A}$ is some invertible matrix; as input only the action of the matrix $\mathbb{A}$ on a vector needs to be supplied. For the present problem this can be achieved with cost $O(D^3)$, using a contraction scheme along the lines of Appendix B.

Achieving $O(D^3)$ scaling assumes that, in the limit of large bond dimension D, the number of iterations needed for the biconjugate gradient algorithm to converge to some desired accuracy scales as roughly a constant. In practice this may not always be the case. Using the biconjugate algorithm as described above, $O(D^3)$ scaling is indeed spoiled for a general uCPS TDVP computation. This is ultimately related to the fact that uCPS, viewed as a restriction of uMPS, fixes nearly all the gauge freedom, which in general causes G to become badly conditioned in the course of a TDVP run - even when Schmidt values of the uCPS state itself are much larger than machine precision. One solution to this problem is to make a judicious choice of a pre-conditioner matrix. In general, this refers to a matrix $\mathbb{M}^{-1}$, in $\mathbb{M}^{-1}\mathbb{A}\vec{x} = \mathbb{M}^{-1}\vec{b}$, which can be employed in the iterative algorithm in order to make the problem better conditioned. For the present case, taking $\mathbb{M}^{-1} = \tilde{G}$ in fact seems to be the optimal choice (again, the iterative algorithm only needs to be supplied with the action of $\tilde{G}$ on a vector, which can be achieved with cost $O(D^3)$). It should be noted that without the pre-conditioning step, for most of the examples in this paper, it would only have been practically feasible to compute quenches and ground state approximations with very small overlap sizes, and many of the computations performed in this paper would not have been accesible.

[1] R. Darradi, O. Derzhko, R. Zinke, J. Schulenburg, S. Krüger, and J. Richter, Physical Review B **78**, 214415 (2008).

[2] E. Altman and A. Auerbach, Physical Review B **65**, 104508 (2002).

[3] L. Isaev, G. Ortiz, and J. Dukelsky, Physical Review B **79**, 024409 (2009).

[4] H. J. Changlani, J. M. Kinder, C. J. Umrigar, and G. K.-L. Chan, Phys. Rev. B **80**, 245116 (2009), arXiv:0907.4646 [cond-mat.str-el].

[5] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, New Journal of Physics **11**, 083026 (2009), arXiv:0905.3898 [cond-mat.str-el].

[6] F. Mezzacapo and J. I. Cirac, New Journal of Physics **12**, 103039 (2010), arXiv:1006.4480 [cond-mat.str-el].

[7] K. H. Marti, B. Bauer, M. Reiher, M. Troyer, and F. Verstraete, New Journal of Physics **12**, 103008 (2010), arXiv:1004.5303 [physics.chem-ph].

[8] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, New Journal of Physics **11**, 083026 (2009).

[9] F. Mezzacapo and J. I. Cirac, New Journal of Physics **12**, 103039 (2010).

[10] S. Al-Assam, S. Clark, C. Foot, and D. Jaksch, Physical Review B **84**, 205108 (2011).

[11] E. Neuscamman and G. K.-L. Chan, Physical Review B **86**, 064402 (2012).

[12] T. Đurić, N. Chancellor, and I. F. Herbut, Physical Review B **89**, 165123 (2014).

[13] T. Đurić, N. Chancellor, P. J. Crowley, P. Di Cintio, and A. G. Green, Physical Review B **93**, 085143 (2016).

[14] M. Fannes, B. Nachtergaele, and R. Werner, Comm. Math. Phys. **144**, 443 (1992).

[15] F. Verstraete, V. Murg, and J. Cirac, Advances in Physics **57**, 143 (2008).

[16] F. Verstraete and J. I. Cirac, eprint arXiv:cond-mat/0407066 (2004), cond-mat/0407066.

[17] G. Vidal, Physical Review Letters **99**, 220405 (2007), arXiv:cond-mat/0512165.

[18] G. Vidal, Physical Review Letters **101**, 110501 (2008), arXiv:quant-ph/0610099.

[19] J. Haegeman, T. Osborne, H. Verschelde, and F. Verstraete, Physical Review Letters **110**, 100402 (2013).

[20] M. Hastings, Journal of Statistical Mechanics: Theory and Experiment **2007**, P08024 (2007).

[21] D. Gottesman and M. Hastings, New journal of physics **12**, 025002 (2010).

[22] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Physical Review Letters **98**, 140506 (2007), quant-ph/0611050.

[23] D. Gottesman, (1997), arXiv:quant-ph/9705052 [quant-ph].

[24] A. Yu. Kitaev, Annals Phys. **303**, 2 (2003), arXiv:quant-ph/9707021 [quant-ph].

[25] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Physical Review Letters **100**, 040501 (2008), arXiv:0708.1567 [quant-ph].

[26] H. J. Changlani, J. M. Kinder, C. J. Umrigar, and G. K.-L. Chan, Physical Review B **80** (2009), 10.1103/PhysRevB.80.245116.

[27] E. Neuscamman, H. J. Changlani, J. M. Kinder, and G. K.-L. Chan, Physical Review B **84** (2011), 10.1103/PhysRevB.84.205132.

[28] J. Haegeman, T. Osborne, and F. Verstraete, Phys. Rev. B **88**, 075133 (2013), arXiv:1305.1894 [quant-ph].

[29] I. P. McCulloch, ArXiv e-prints (2008), arXiv:0804.2509 [cond-mat.str-el].

[30] J. Haegeman, J. Cirac, T. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Physical Review Letters **107**, 070601 (2011).

[31] L. Tagliacozzo, T. de Oliveira, S. Iblisdir, and J. Latorre, Phys. Rev. B **78**, 024410 (2008), arXiv:0712.1976 [cond-mat.stat-mech].

[32] Y. Fuji, F. Pollmann, and M. Oshikawa, Physical Review Letters **114**, 177204 (2015), arXiv:1409.8616 [cond-

mat.str-el].
[33] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, ArXiv e-prints (2014), arXiv:1408.5056 [quant-ph].
[34] F. Pollmann, S. Mukerjee, A. Turner, and J. Moore, Physical Review Letters **102**, 255701 (2009), arXiv:0812.2903 [cond-mat.str-el].
[35] B. Pirvu, G. Vidal, F. Verstraete, and L. Tagliacozzo, Physical Review B **86**, 075117 (2012), arXiv:1204.3934.
[36] V. Stojevic, J. Haegeman, I. P. McCulloch, L. Tagliacozzo, and F. Verstraete, Phys. Rev. B **91**, 035120 (2015), arXiv:1401.7654 [quant-ph].
[37] C. Holzhey, F. Larsen, and F. Wilczek, Nuclear Physics B **424**, 443 (1994), arXiv:hep-th/9403108.
[38] P. Calabrese and J. Cardy, Journal of Statistical Mechanics: Theory and Experiment **2004**, P06002 (2004).
[39] A. Silva, Phys. Rev. Lett. **101**, 120603 (2008).
[40] F. Andraschko and J. Sirker, Phys. Rev. B **89**, 125120 (2014), arXiv:1312.4165 [cond-mat.str-el].
[41] M. Heyl, A. Polkovnikov, and S. Kehrein, Physical Review Letters **110**, 135704 (2013), arXiv:1206.2505 [cond-mat.stat-mech].
[42] C. Karrasch and D. Schuricht, Phys. Rev. B **87**, 195104 (2013), arXiv:1302.3893 [cond-mat.str-el].
[43] J. Haegeman, M. Mariën, T. J. Osborne, and F. Verstraete, ArXiv e-prints (2012), arXiv:1210.7710 [quant-ph].